AD-A146 209 DISTRIBUTED SENSOR NETWORKS(U) MASSACHUSETTS INST OF TECH LEXINGTON LINCOLN LAB R T LACOSS 30 SEP 83 ESD-TR-84-002 F19628-80-C-0002 1/1

UNCLASSIFIED F/G 9/2 NL

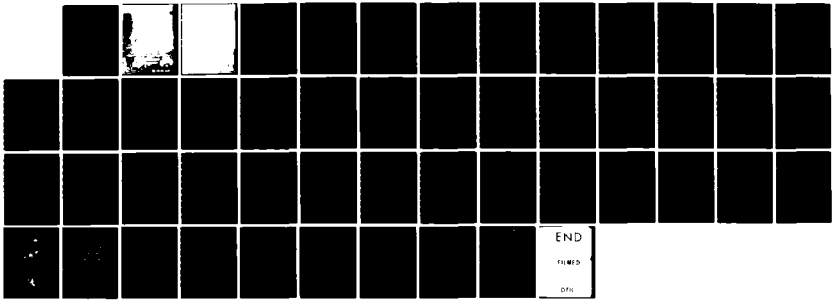

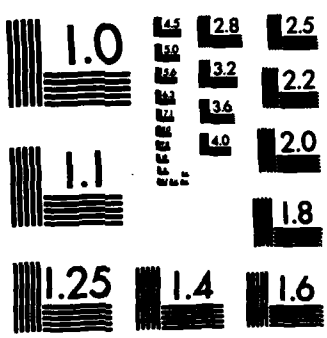

MICROCOPY RESOLUTION TEST CHART

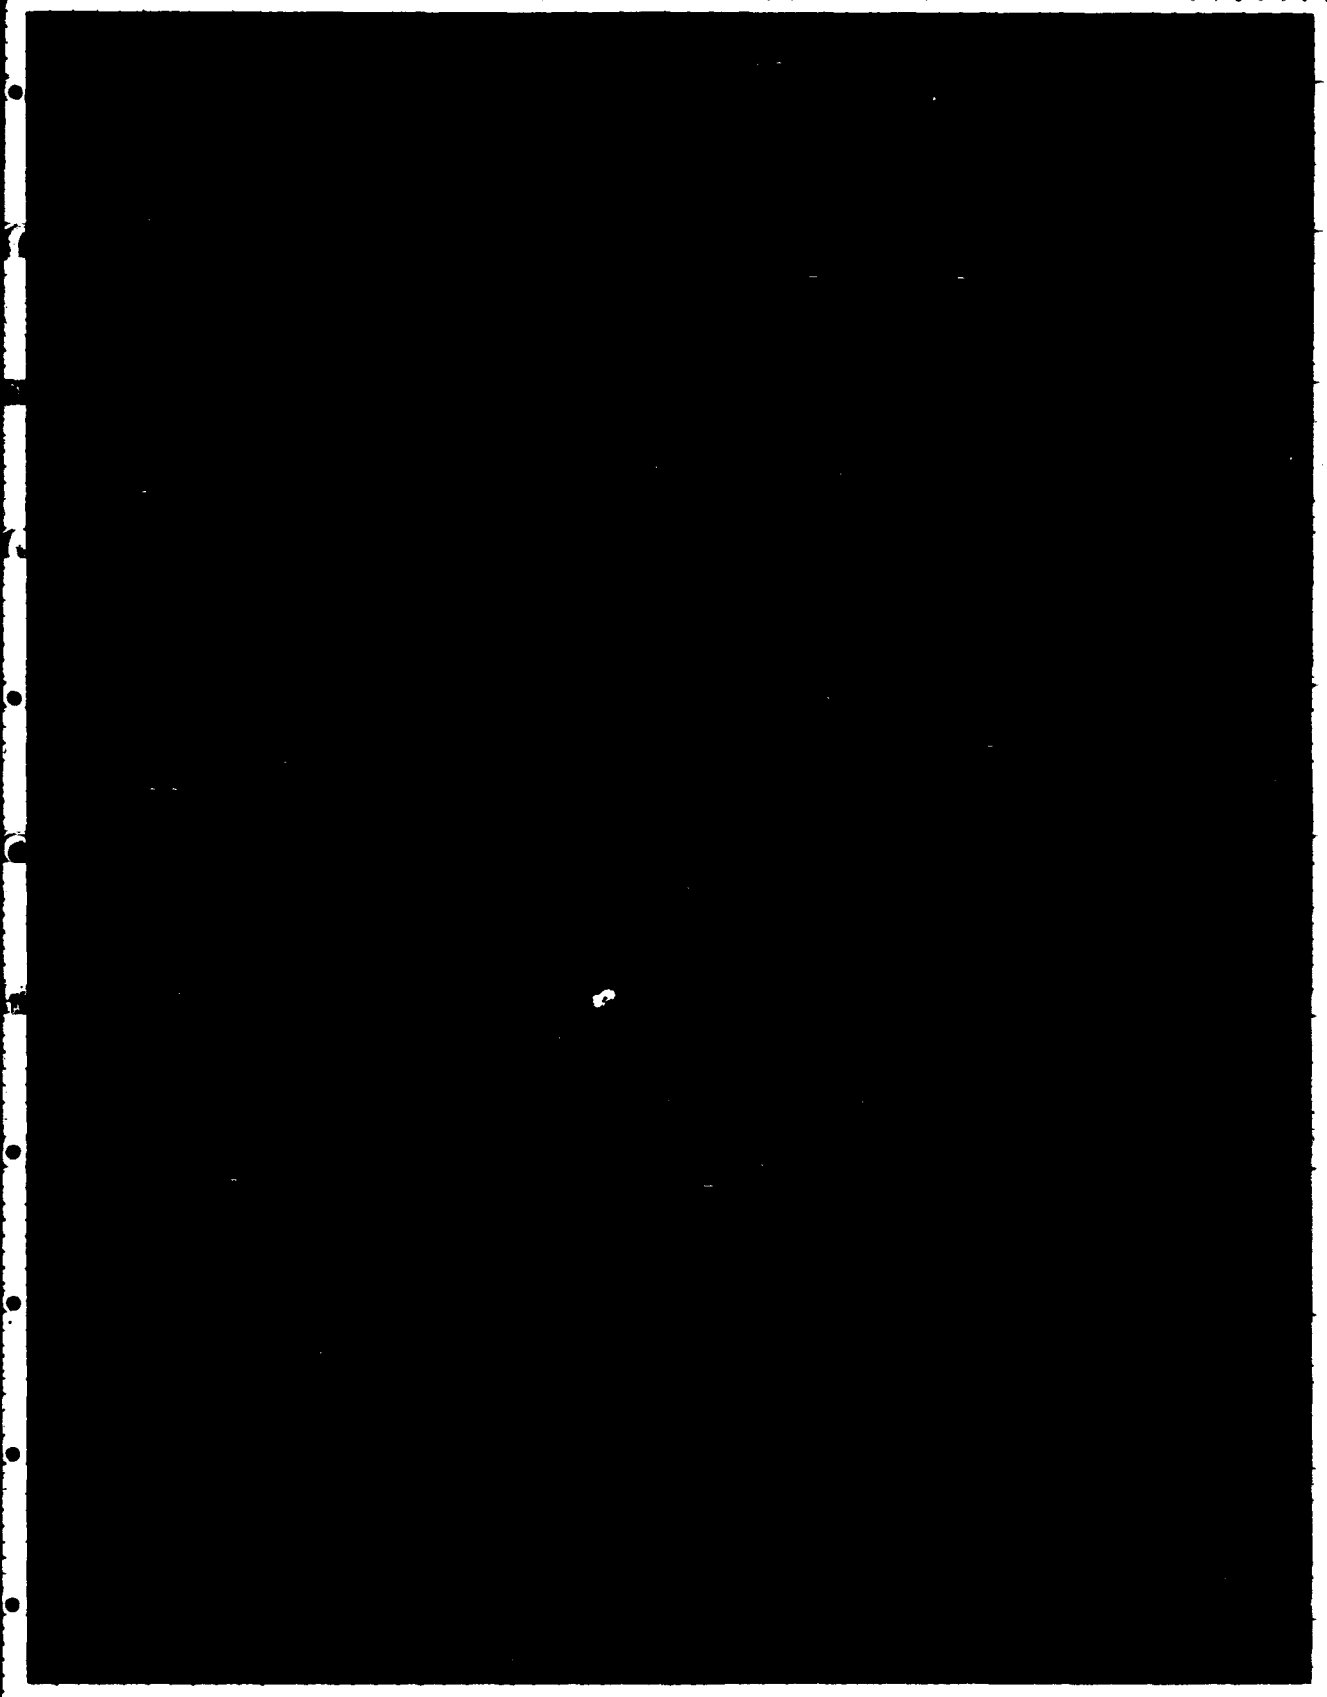

MASSACHUSETTS INSTITUTE OF TECHNOLOGY
LINCOLN LABORATORY

DISTRIBUTED SENSOR NETWORKS

SEMIANNUAL TECHNICAL SUMMARY REPORT
TO THE
DEFENSE ADVANCED RESEARCH PROJECTS AGENCY

1 APRIL — 30 SEPTEMBER 1983

ISSUED 12 JUNE 1984

LEXINGTON MASSACHUSETTS

ABSTRACT

This report describes the work performed on the DARPA Distributed Sensor Networks Program at Lincoln Laboratory during the period 1 April through 30 September 1983.

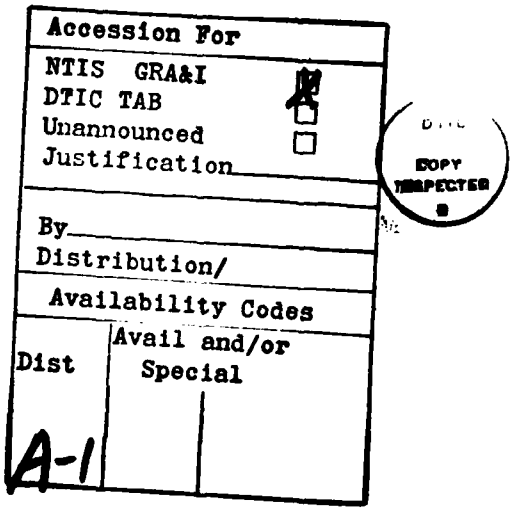

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

DISTRIBUTED SENSOR NETWORKS

I. INTRODUCTION AND SUMMARY

The Distributed Sensor Networks (DSN) program is aimed at developing and extending target surveillance and tracking technology in systems that employ multiple spatially distributed sensors and processing resources. Such a system would be made up of sensors, data bases, and processors distributed throughout an area and interconnected by an appropriate digital data communication system. The detection, tracking and classification of low-flying aircraft has been selected to develop and evaluate DSN concepts in the light of a specific system problem. A DSN test bed has been developed and is being used to test and demonstrate DSN techniques and technology. The sensors presently in use are small arrays of microphones. The overall concept calls for a mix of sensor types. Visible TV sensors are scheduled for integration into the test bed in the near future. This Semiannual Technical Summary (SATS) reports results for the period 1 April through 30 September 1983.

A new wideband acoustic array processing algorithm has been developed that promises to provide better detection and azimuth determination performance than previous algorithms with less computational load. The new method is based upon analysis of the instantaneous intensity pattern in an acoustic array plane rather than upon standard delay-and-sum array processing concepts. One major advantage of the new method is that it can make use of the entire frequency band of broadband sources while requiring no more processing than standard methods require to perform spatial analysis at a single temporal frequency. Initial tests of the new algorithm with experimental acoustic data have indicated that it will perform better than previously available methods. Section II provides more details about this algorithm and the initial experimental results.

Error analysis studies for the possible position and reflection methods of target location using two acoustic arrays have been completed and the conclusions are summarized in Section III. It was judged that the algorithms provide comparable performance with the exception that a significant disadvantage of the reflection method is that it must estimate signal emission times as well as target locations. However, a tentative decision to convert the test bed from the reflection method to the possible position method is moot because we have now identified a superior distributed tracking approach that will involve neither of those algorithms.

Section IV reports upon significant progress that has been made during this reporting period in the development of new distributed tracking algorithms based upon state-of-the-art distributed estimation techniques. Extended Kalman filters were developed to perform real-time acoustic target tracking with azimuth measurements as inputs. A new real-time track initiation algorithm was developed. A distributed version of the acoustic tracking algorithm

was developed that uses Bayesian methods for track combining and works in conjunction with an adaptive communication policy to reduce internodal communication. Simulations were done to test the algorithms for single targets, multiple targets and maneuvering targets. This work has established that there are no fundamental problems with the new distributed tracking approach and that it should be further developed and used in the DSN test bed.

Section V summarizes some exploratory investigations that have been started with the goal of eventually using AI methods to improve DSN system performance.

As reported in Section VI, additional data acquisition experiments were conducted at Hanscom Air Force Base to obtain one- and two-target data for signal processing and tracking algorithm evaluation.

Section VII summarizes progress made with improvements to the test-bed system. These include integration of new standard nodal computers, radios and a user workstation. System software for the new nodal configuration has been completed and tracking software has been converted to the new configuration. Measurements were made of interprocessor message communication rates within a node and they appear to be adequate to meet our test-bed requirements.

II. WIDEBAND SENSOR ARRAY PROCESSING METHOD

Through a fundamentally different approach to sensor array processing, we have developed a new acoustic bearing estimation algorithm which offers significant advantages over previously available algorithms. In contrast to traditional methods, this algorithm makes simultaneous use of target energy over a band of frequencies without excessive computational requirements. This is particularly useful in situations involving jet data, multiple aircraft or strong interference noise. Preliminary evaluation of the algorithm indicates improved performance over previous methods while reducing computational load. The following reviews the motivation for developing a new method, presents a qualitative description of the new method and initial processing results. A more detailed technical treatment can be found in "Direction Determination of Wideband Signals".*

Conventional sensor array processing techniques are variants of delay-and-sum processing followed by a power detector (see Figure II-1). For each direction of arrival, there is a corresponding set of sensor time delays which result in plane wave signals from the selected

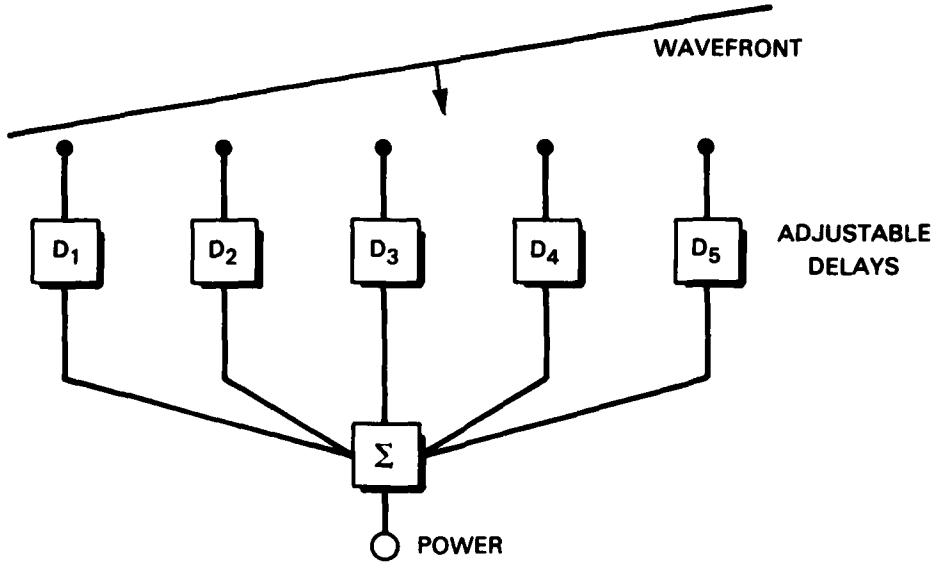

133841-N-01

Figure II-1. Delay-and-sum processing.

*S.H. Nawab, F.U. Dowla and R.T. Lacoss, "Direction Determination of Wideband Signals," Submitted to IEEE Trans. on ASSP, August 1983.

direction being added coherently. Direction estimation is done by searching through all possible directions to find directions that produce power peaks. If the planewave of interest is monochromatic with frequency f, then the sensor delays can be replaced by phase shifts. The overall process then becomes equivalent to Bartlett estimation of the frequency-wavenumber spectrum. Substituting any high-resolution spectral estimation technique for Bartlett estimation yields other modern variants of sensor array processing, including the MLM method that has been used in the DSN test bed*. For our present purposes, the most relevant feature of these techniques is that direction estimation is done one frequency at a time. Consequently, the computational requirements of these methods increase linearly with the number of target frequencies to be considered.

Acoustic energy received at a DSN node is often spread over a wide band of frequencies. Sources such as helicopters are periodic with many harmonics. Jet aircraft tend to produce less structured peaks and have energy more uniformly spread over a broad bandwidth. When targets have such wideband energy and their temporal spectra can be estimated in advance, one approach in traditional array processing is to determine direction for the frequency with highest spectral power. However, choosing such a frequency does not guarantee a good signal-to-noise ratio; there may be greater interfering energy for that frequency than in some other portions of the spectrum. The selection of analysis frequencies is further complicated when there are multiple targets since only the most dominant target may be detected. This leads to the need to analyze a large number of frequencies. In addition, algorithms are needed to combine data from different frequencies when broadband signal-to-noise ratios are large even if the signal-to-noise ratios are small for single frequencies.

The new algorithm solves these problems. The approach focuses on the notion of measuring instantaneous intensity patterns in the array plan rather than on delay-and-sum concepts. Consider a plane wave impinging on a sensor array as illustrated in Figure II-2. Since at any instant the signal intensity is constant in a plane orthogonal to the direction of propagation, the intensity pattern on the array is constant along lines of intersection between those planes and the plane of the array. As shown in Figure II-2, these lines of constant amplitude are perpendicular to the direction along the dashed line, OP, which denotes the bearing corresponding to the angle Θ_{AZ}. Furthermore, the variation along OP is directly proportional to the temporal variations in the planewave. It then follows that the two-dimensional Fourier transform of the intensity pattern in the array plane has the form illustrated in Figure II-3. The orientation of the 'ridge' in this Fourier transform corresponds to the azimuth angle Θ_{AZ} of the plane wave. Furthermore, the variation of the power along the ridge is proportional to the temporal spectrum of the corresponding plane wave.

---

*J. Capon, "High-Resolution Frequency-Wavenumber Spectrum Analysis," Proc. IEEE, 57, 1408 (1969) DDC AD-6968800.

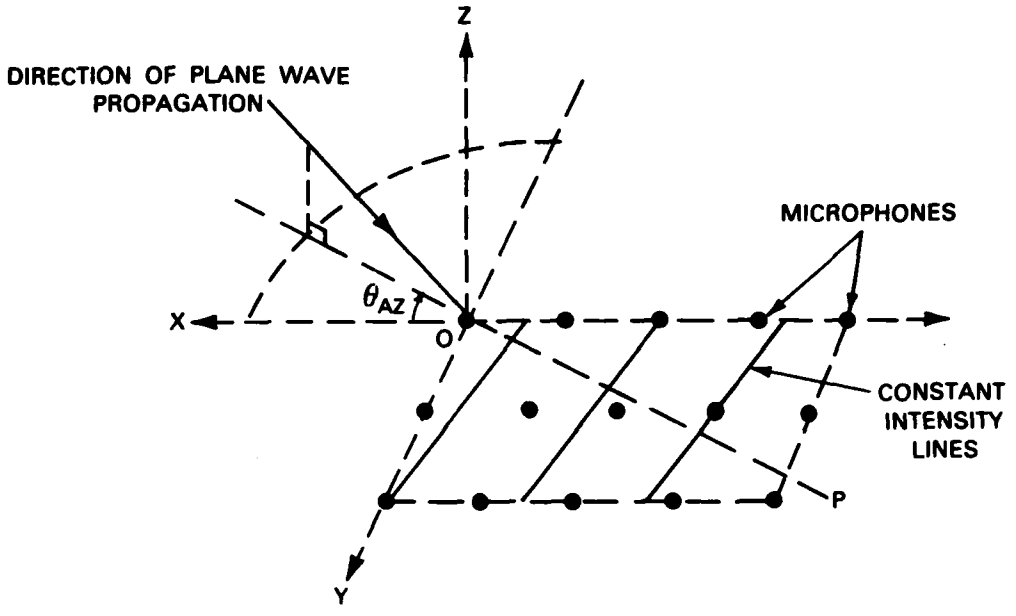

133843-N-02

Figure II-2. Plane wave impinging on sensor array.

Note that the transform illustrated in Figure II-3 does not allow us to distinguish between azimuths Θ_{AZ} and Θ_{AZ} + 180°. This is because temporal power spectra of real signals are symmetric about the origin. We have solved this problem by incorporating in our algorithm an efficient procedure for computing samples of the complex analytic representation of the sensor signals. This results in the suppression of the negative frequency band and thus allows unique determination of azimuth angle or bearing.

In the above discussion, many of the technical details of the algorithm have been glossed over. For example, instead of measuring an intensity pattern, zero-delay covariances of the sensor signals are measured and instead of the 2-D Fourier transform, the high resolution MLM method is used for 2-D spectral estimation.

A comparison of the computations involved in the new algorithm and the existing DSN nodal algorithm is represented in the block diagram of Figure II-4. Starting from the K time series received at the array sensors, the sequence of operations for the two algorithms is shown on either side of the figure. As indicated at the bottom of the figure, if the outputs of the old method are integrated over all frequencies, then the result is equivalent to the output of the new method. It should be noted that many of the operations are the same in both algorithms. The essential difference is that the old algorithm repeats these operations as many times as there are frequencies to analyze. The new algorithm covers an entire band of frequencies for the same amount of computation as required for a single frequency with the old algorithm.

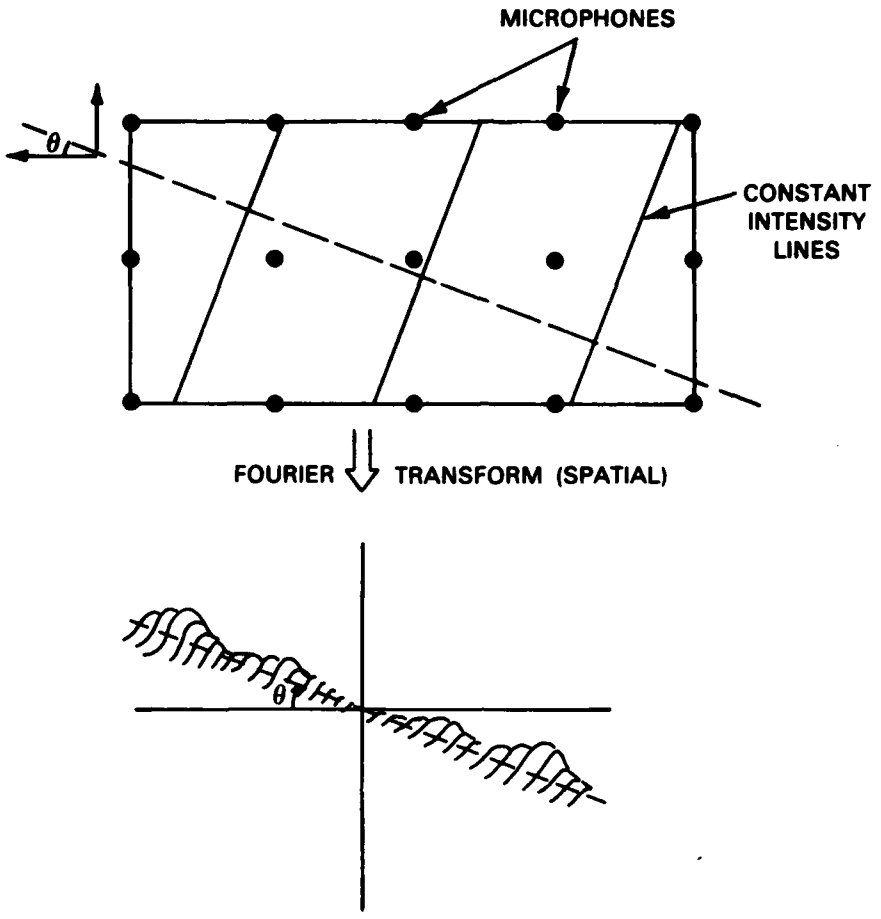

Figure II-3. Fourier transform of intensity pattern in array plane.

To demonstrate the capabilities of the new algorithm, we have tested it with experimental acoustic DSN data. The data was recorded at a 2048-Hz sampling rate on a microphone array with the configuration of Figure II-5. In the following description of our results, north corresponds to a bearing of 0 degrees.

The first example is for a UH-1 helicopter approaching the array from the east, a distance three kilometers away. Figure II-6(a) is a contour plot of the output of the new algorithm. The plot clearly shows dominant signal power along the 96-degree bearing, which corresponds to the eastern bearing of the UH-1. For this same output, the dotted line in Figure II-6(b) is the radial cross section at the 96-degree bearing. Also shown there is an estimate of the average temporal power spectrum for the nine sensor signals. The radial cross section is clearly a smoothed estimate of the temporal spectrum as predicted by the theory. The calculations for obtaining Figure II-6 used two seconds of data and the frequency band from 0 to 128 Hz.

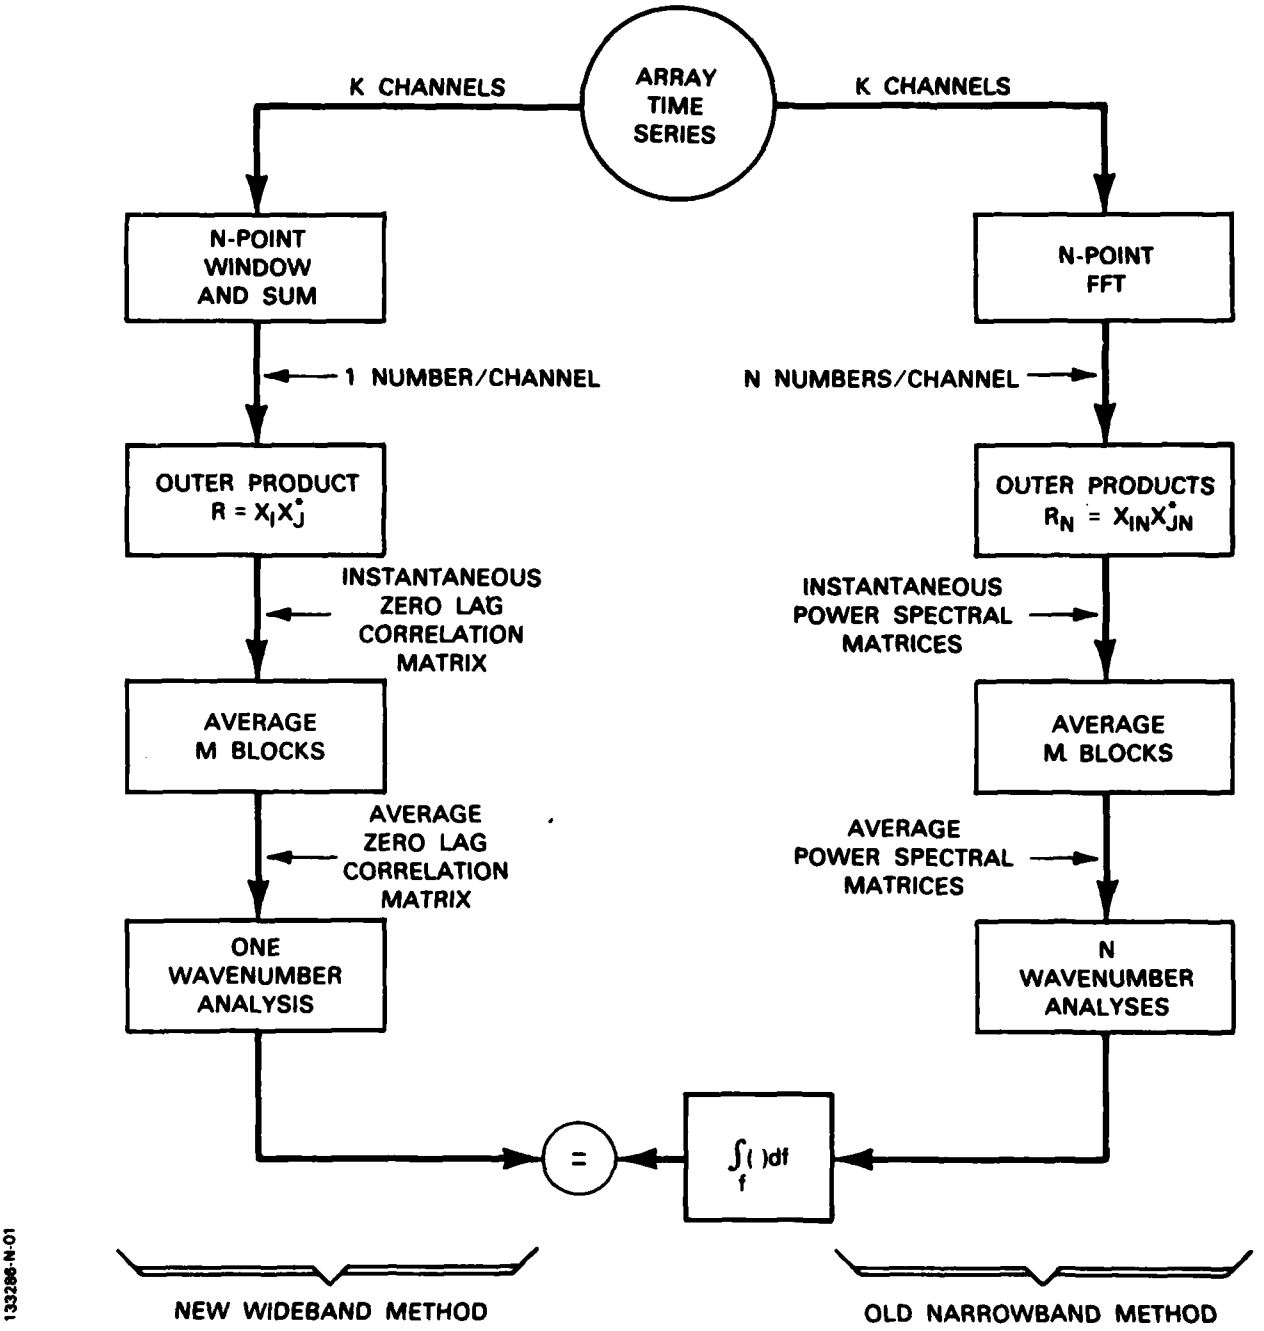

Figure II-4. Computational comparison of old and new algorithms.

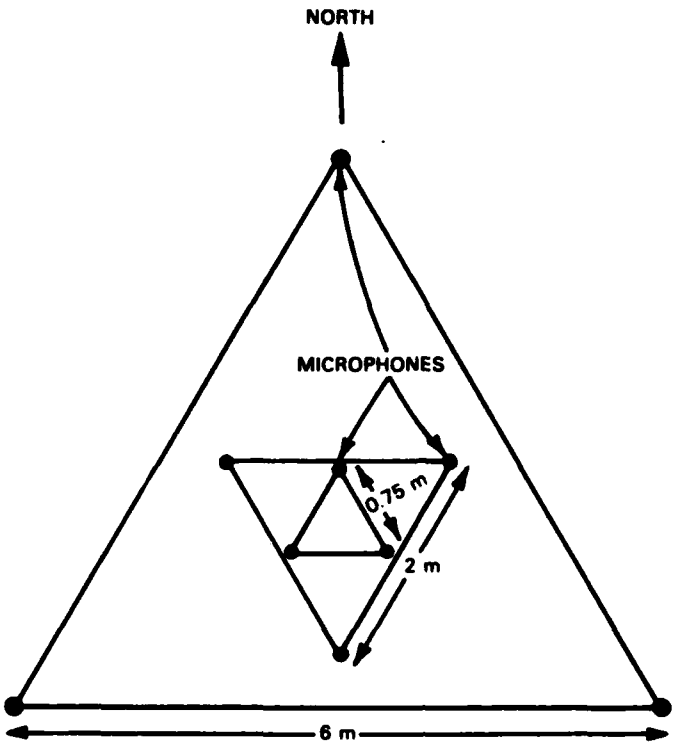

Figure II-6. Nine-sensor DSN array.

Figure II-7 illustrates another way of viewing the output of the new algorithm. The results shown are for data that included signals from a UH-1 helicopter and a propeller fixed-wing aircraft. The two-dimensional power output of the algorithm has been converted to a one-dimensional representation by radially integrating the two-dimensional output for each azimuth. The result is broadband integrated power as a function of azimuth. Figure II-7 shows three such plots for three analysis intervals, separated by five seconds from each other. During this time, the helicopter stayed at 90-degree azimuth while the propeller aircraft went from a 120-degree azimuth, through the 90-degree azimuth, and finally reached the 60-degree azimuth.

Our final example compares the output of the new algorithm with those of the old one. In this case, we selected the highest azimuthal peak detected for each analysis interval and plotted the location of these peaks as a function of time. The data were for a single UH-1 helicopter flying past an array on a straight line trajectory. The results for the old algorithm are shown in Figure II-8(a). This analysis was done every two seconds for the eight spectral peak frequencies with highest power. The results of processing with the new

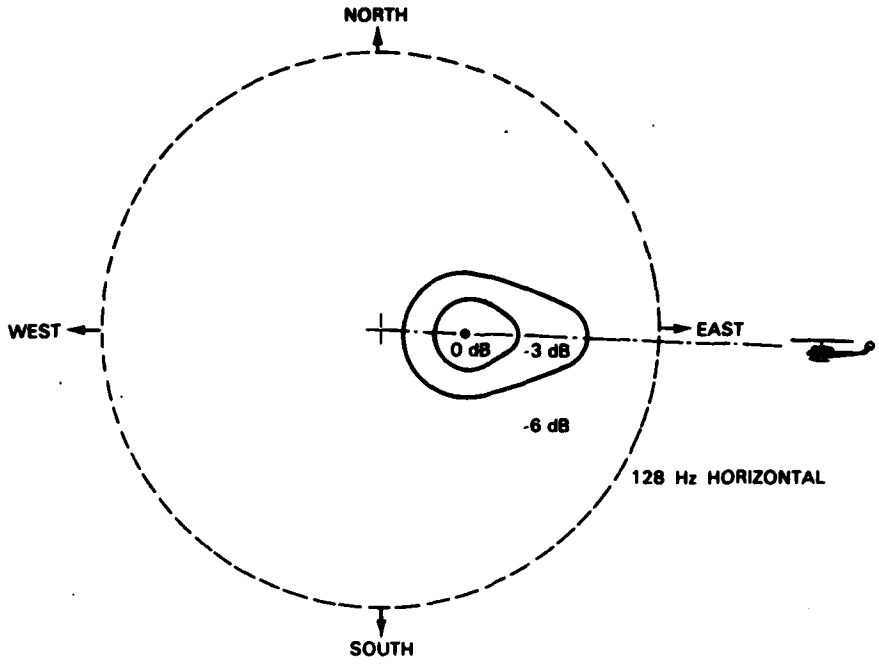

133845-N-02

Figure II-6(a). Contour plot of directional spectrum for helicopter with easterly bearing.

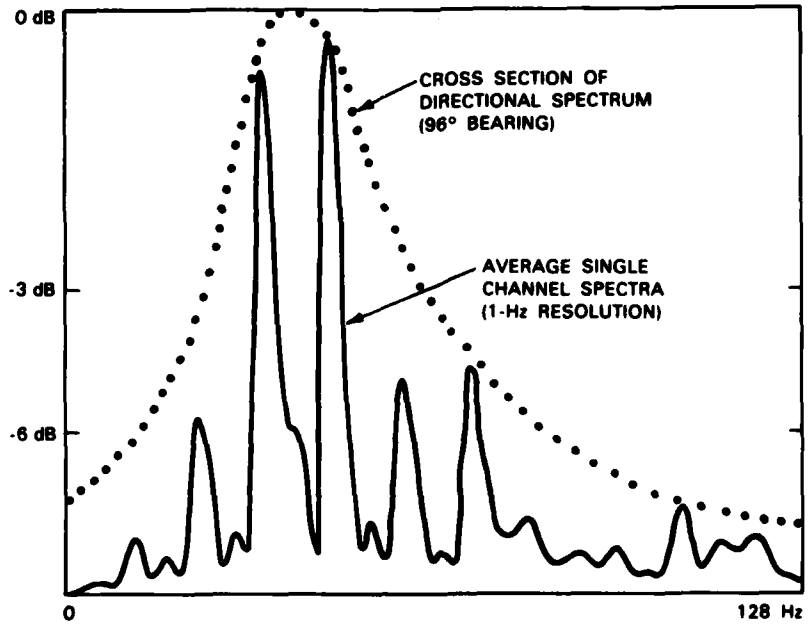

133846-N-01

Figure II-6(b). Radial cross section at 96-degree bearing for spectrum in Figure II-6a.

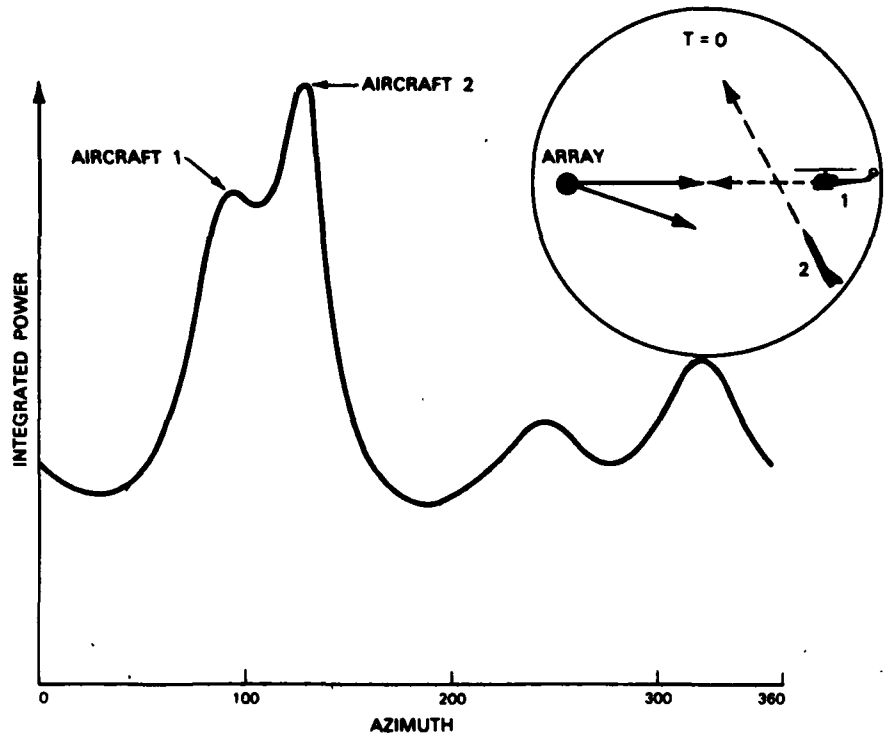

Figure II-7(a). Radially integrated target power and target configuration at T = 0.

algorithm are given in Figure II-8(b). The superior performance of the new algorithm is self-evident from the figure. The output from the new algorithm does not exhibit the degree of random fluctuations found in the output of the old algorithm.

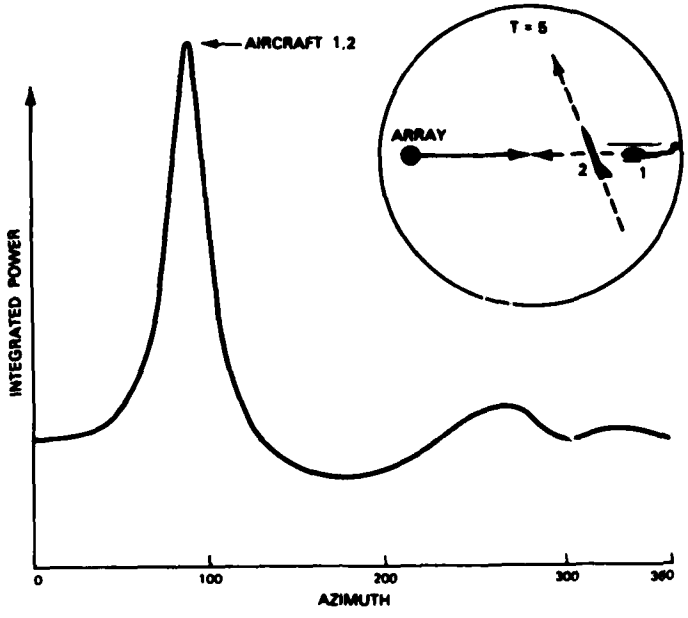

133948-N-02

Figure II-7(b). Radially integrated target power and target configuration at T = 5.

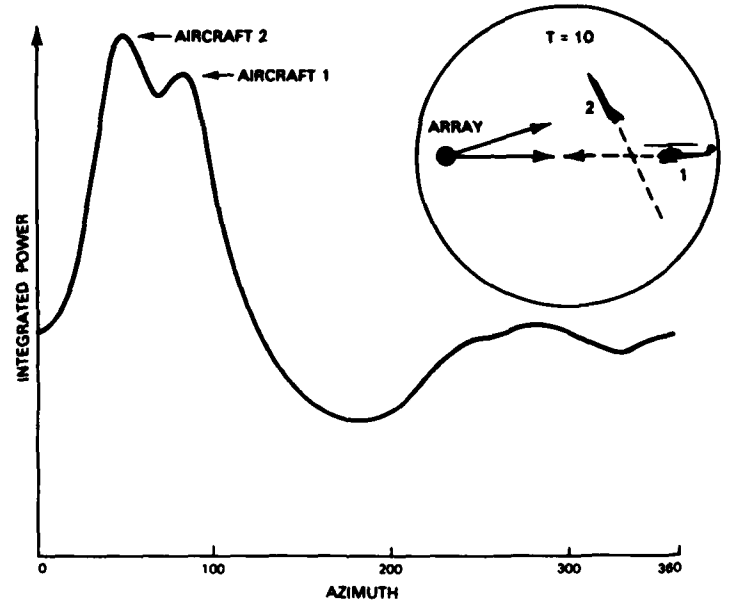

133949-N-02

Figure II-7(c). Radially integrated target power and target configuration at T = 10.

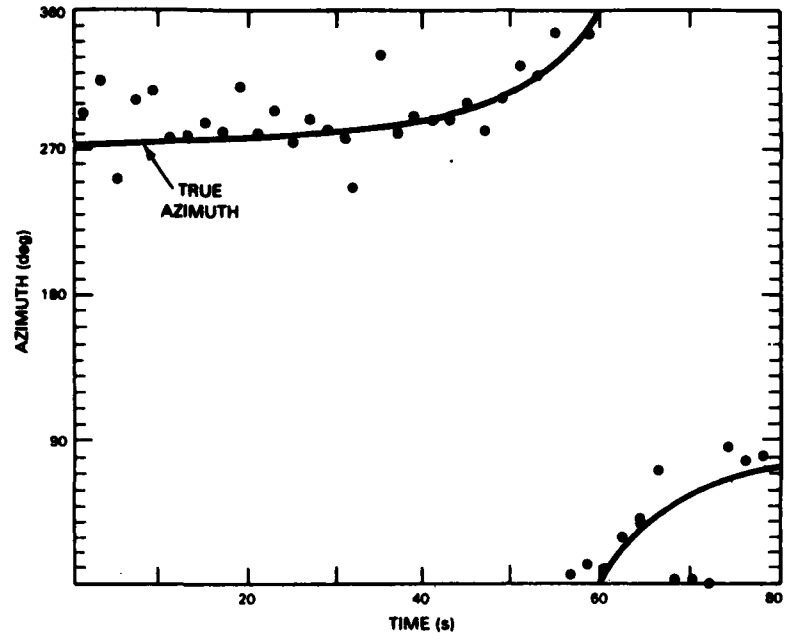

137024-R

Figure II-8(a). Azimuth estimates as a function of time with old algorithm. Solid line represents ground truth.

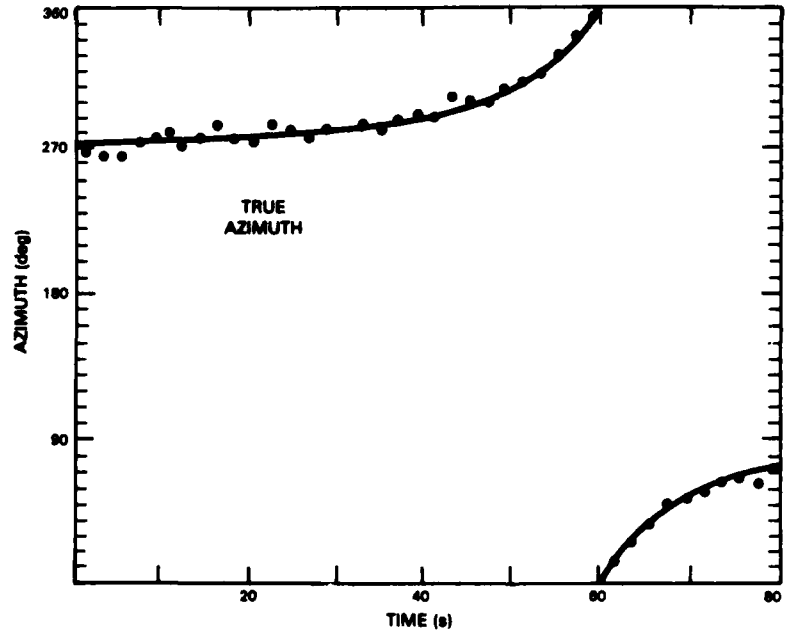

137023-R

Figure II-8(b). Azimuth estimates as a function of time with new algorithm. Solid line represents ground truth.

III. ERROR ANALYSIS FOR TWO-NODE TARGET LOCATION ALGORITHMS

Initial analytical results dealing with the sensitivity of estimated target locations to errors in acoustic azimuth measurements and to errors in nodal locations, were presented in a recent SATS*. The analysis was carried out for two different methods of calculating target locations, using azimuth estimates from two acoustic arrays. The two methods are the possible position method† and the reflection method‡. Given target time, the possible position method determines target locations by converting azimuth tracks from each node into target positions and locating the target where the positions from both nodes coincide. The reflection method operates by converting an azimuth measurement from one node into predicted azimuth measurements at a second node and calculates a target location when a prediction matches the azimuth track at the second node. The reflection method is the algorithm now used in the DSN test bed, and we were considering replacing it with the possible position method. The analysis was continued into this reporting period with emphasis upon the role of spatial diversity in a DSN, the changing utility of different sensor pairs along a target track, and the sensitivity of estimated target locations to errors in sensor locations, as well as errors in target azimuth measurements. Following is a summary of the overall conclusions and examples of the more recent sensitivity calculations that were used to reinforce and extend the conclusions of the initial analysis.

The objective of the analysis was to decide whether the reflection method should be replaced by the possible position method in the test bed. The general conclusion was affirmative but, as summarized in Section IV, subsequent research into distributed tracking options has resulted in a new tracking approach which offers several advantages over these previously developed algorithms.

The general conclusions concerning the possible position and reflection methods of target location were:

(1) Target location error sensitivities are comparable for the two algorithms.

(2) Target location errors due to azimuth measurement errors will be more significant than those due to nodal location errors for nodal location errors of a few hundred meters or less.

---

*Distributed Sensor Networks Semiannual Technical Summary, Lincoln Laboratory, M.I.T. (30 September 1982) p. 3.

†Distributed Sensor Networks Semiannual Technical Summary, Lincoln Laboratory, M.I.T. (31 March 1978) p. 40.

‡Distributed Sensor Networks Semiannual Technical Summary, Lincoln Laboratory, M.I.T. (31 March 1980) p. 6.

(3) The usefulness of a node pair changes as a function of time as a target moves along its trajectory, and practical DSN algorithms must take advantage of this fact.

(4) A significant disadvantage of the reflection method is that it estimates acoustic signal emission times which in turn are used to estimate target velocities, and this can result in very large errors in estimated velocities.

Figures III-1 and III-2 show examples of the analytical calculations used to support these conciusions. The calculations are for a target flying a constant velocity straight line trajectory through an L-shaped group of three sensors. This simple scenario makes explicit the changing performance for different node pairs as a target flies through the coverage of several sensors. Figure III-3 shows the geometry of the sensors and the path of the target. Target location error statistics were calculated at regular intervals along the target path for several velocities ranging from 0.1 Mach to 0.9 Mach and for both location algorithms. Figures III-1 and III-2 show calculations for only the reflection method and for 0.1 Mach.

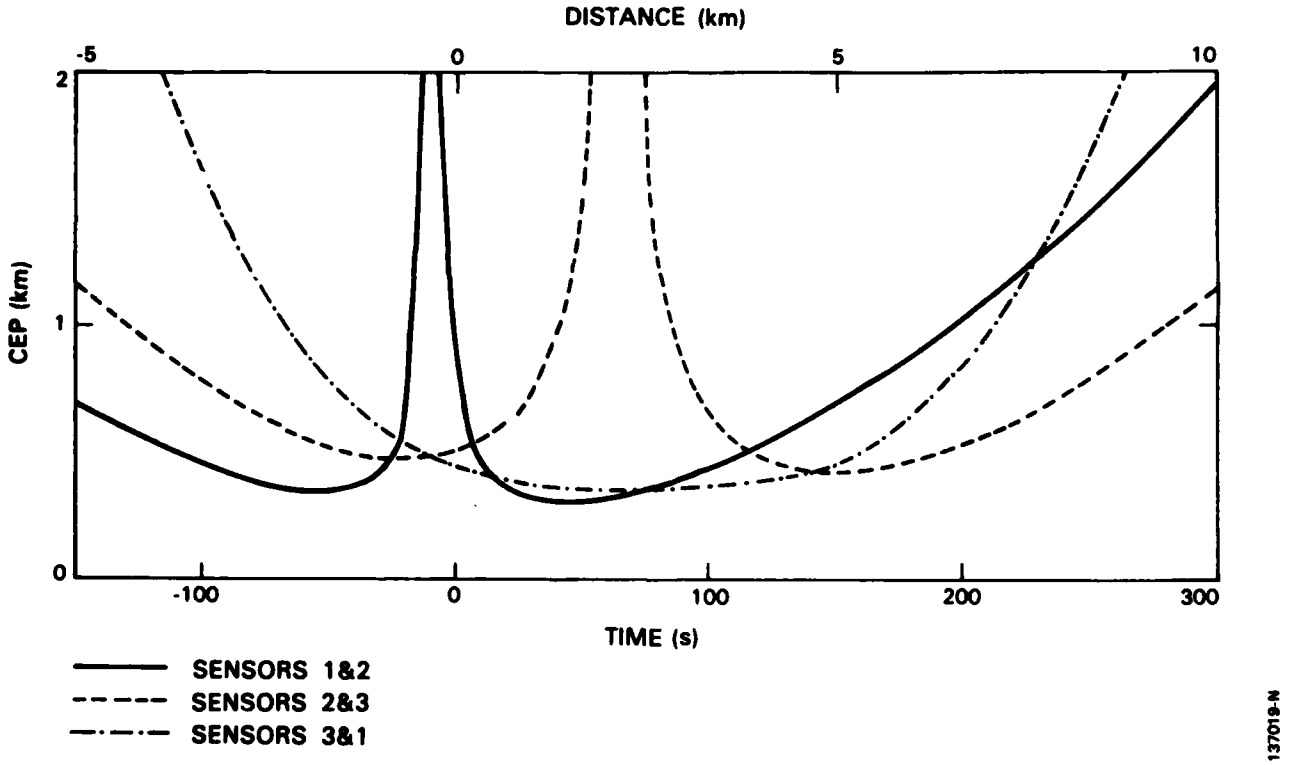

Figure III-1. Location estimate sensitivity of reflection method to azimuth measurement errors: σ = 5°, V_t = 0.1 Mach.

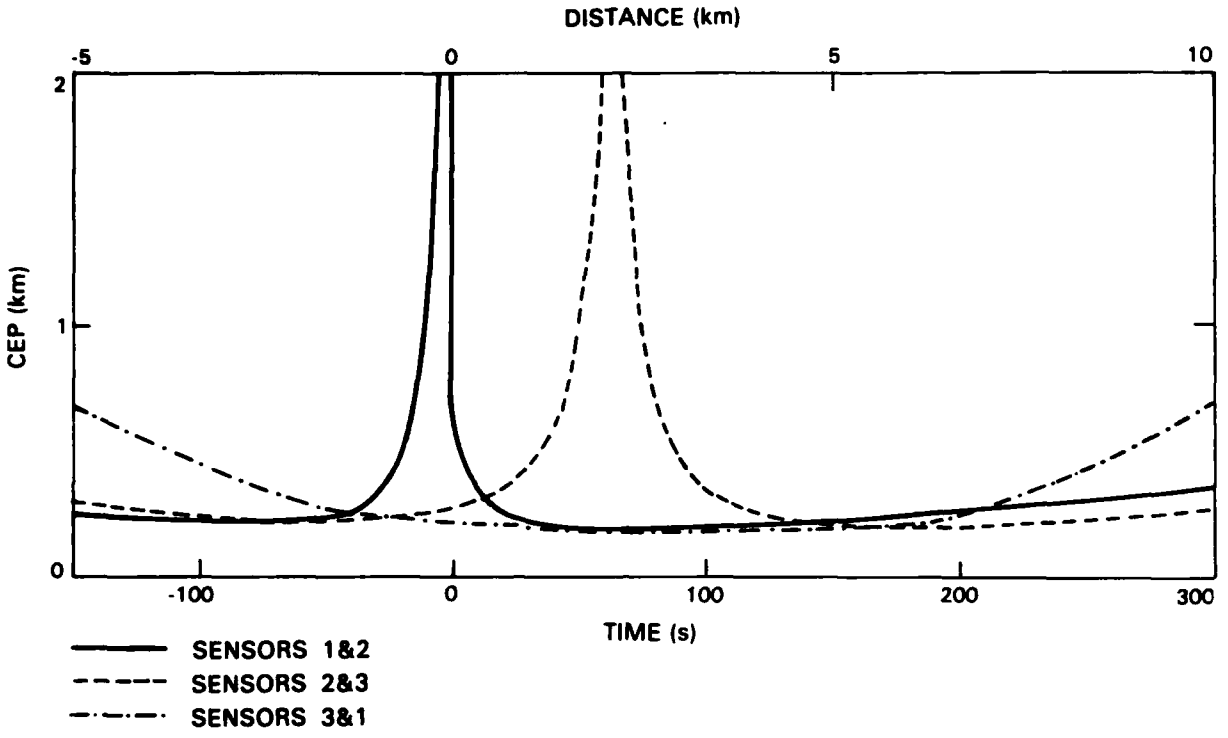

Figure III-2. Location estimate sensitivity of reflection method to sensor location errors: CEP = 200 m, V_t = 0.1 Mach.

Figure III-1 shows the sensitivity of the reflection method to azimuth measurement errors. The azimuth measurement errors were assumed to have a mean of zero and a variance of 5 degrees. Our experience is that this is generally consistent with what we obtain using real acoustic data, although somewhat pessimistic. The zero time in the figure corresponds to the target crossing the line between sensors 1 and 2. The parallel distance scale is relative to the same point and is measured along the target path. The three curves in the figure show the target location circle of error probable (CEP) as a function of time for each pairing of the three sensors. It is clear that at least one pair of sensors produces reasonably accurate target location estimates at each time, but which pair is best varies. Moreover, there are times when one pair or another produces very poor estimates. The same comments are true at different velocities and for the possible position method.

Figure III-2 is similar to III-1 except that it is for 200-meter CEP errors in sensor locations and no errors in azimuth measurement. These nodal location errors are considerably larger than we expect from a DSN self-location system based upon internodal radio ranging measurements, although a detailed analysis of self-location has not yet been completed. There are similar singularities in both figures but, even for these large assumed nodal location errors, the envelope of target location errors is lower for Figure III-2 than for Figure III-1.

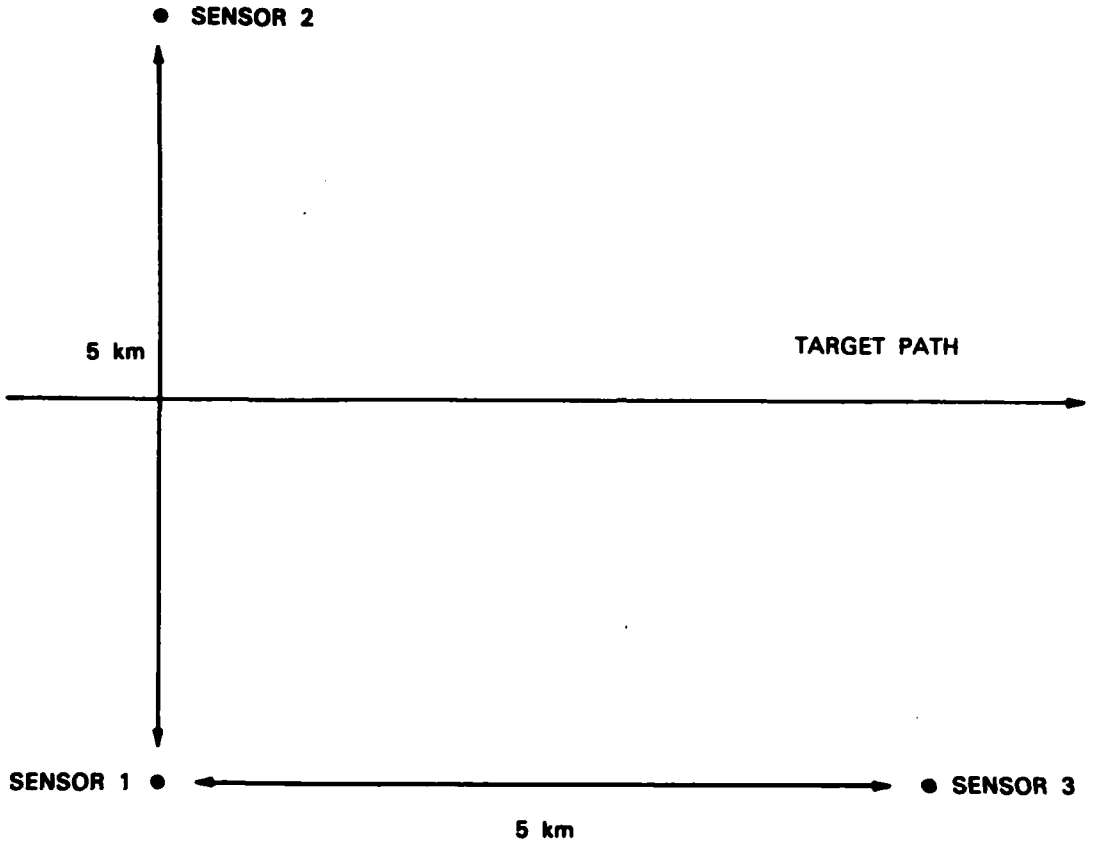

Figure III-3. Sensor locations and target trajectory.

Figure III-4 illustrates the most significant disadvantage of the reflection method of target location. The reflection estimates signal emission times as well as source locations. The figure shows the time estimation errors corresponding to the same scenario as Figure III-1. The quantity plotted is the standard deviation of the temporal estimation error. Again, there are singularities. The envelope of least errors ranges between .5 and 1.5 seconds. This error is significant because target velocities are estimated by:

$$V_{new} + \frac{P_{new} - P_{old}}{T_{new} - T_{old}}$$

where P denotes target position and T denotes time. Locations are estimated whenever new azimuth measurements are obtained and they are now obtained every two seconds. Thus, the errors in the estimates making up the denominator of the above equation can be a significant fraction of the true time difference and very large errors in the velocity estimate can result.

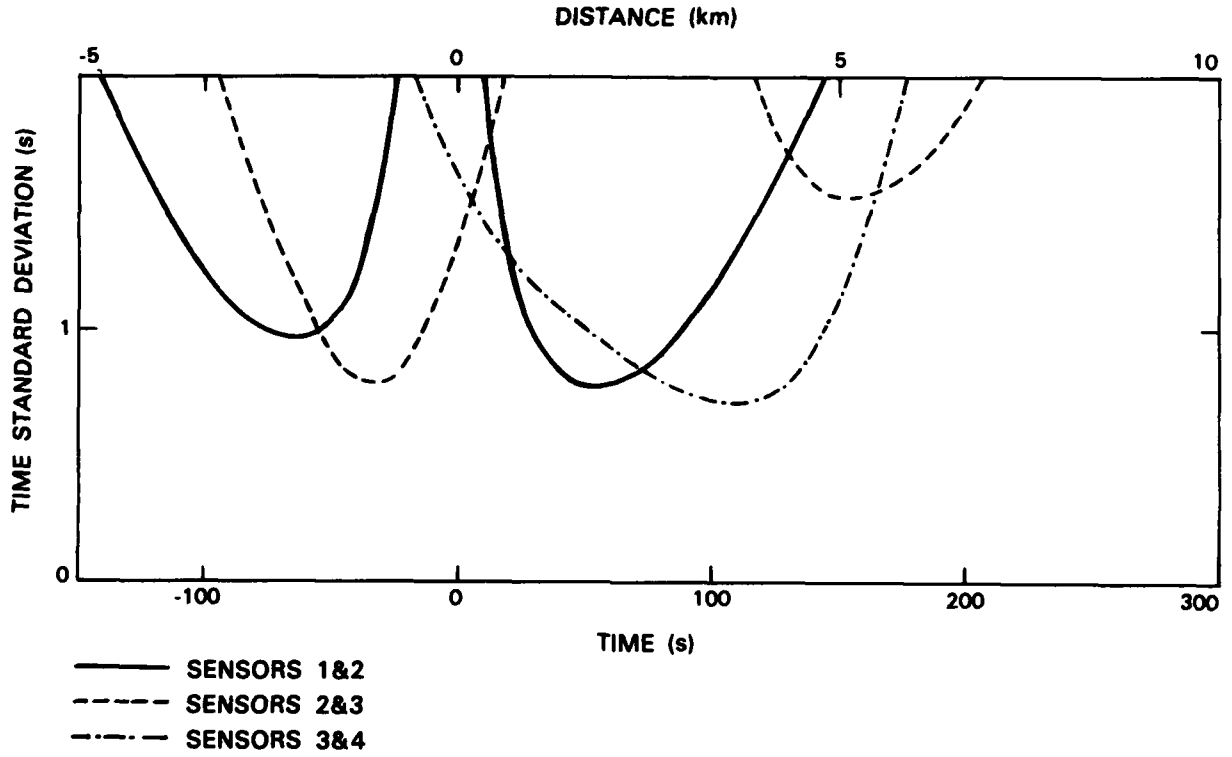

Figure III-4. Time estimate sensitivity of reflection method to azimuth measurement errors: σ = 5°, V_t = 0.1 Mach.

IV. DISTRIBUTED TRACKING ALGORITHMS

An exploratory investigation of new target tracking algorithms based upon state-of-the-art distributed estimation techniques was undertaken during this reporting period. The investigation was motivated by a desire to find algorithms that would provide real-time tracks in spite of acoustic signal propagation delays of many seconds, that would provide tracks using any number of nodes rather than just pairs of nodes, that would optimally take account of spatial diversity in a DSN, and that could operate with a reduced amount of information exchange between nodes. In addition, it appeared desirable to investigate algorithms that could easily handle different kinds of sensors and for which there was a firm theoretical base of understanding.

The focus of this work was on possible fundamental limitations to the application of distributed estimation techniques to DSN acoustic tracking. The goal was to identify roadblocks, if any, resulting from the physical nature of acoustic tracking or to limitations of the techniques. The specific questions we have addressed during this reporting period are:

(1) Can real-time acoustic tracking algorithms be developed that work directly with azimuth measurements as input data?

(2) Can a practical track initiation algorithm be developed that provides real-time target position and velocity estimates in spite of acoustic propagation delays?

(3) Can distributed tracking algorithms be developed that require the exchange of only small amounts of information between nodes?

(4) Can distributed tracking algorithms effectively handle multiple maneuvering targets and false alarms?

We have shown that the answers to all four questions are affirmative and, therefore, that it will be possible to implement an improved DSN tracking system based upon the firm theoretical basis of modern distributed estimation techniques. Following is more detail concerning the investigation of each of the four questions.

A. REAL-TIME CENTRALIZED ACOUSTIC TRACKING ALGORITHMS

The feasibility of updating target tracks directly from acoustic azimuth measurements has been demonstrated by developing and testing a centralized extended Kalman filter [Reference IV-1] algorithm for this task. The algorithm processes acoustic azimuth measurements from any number of nodes. This has established that: (1) tractable linearized equations can be developed to relate the evolution of target states to acoustic azimuth measurements, and that (2) the linearized equations are a sufficiently accurate approximation for a practical tracking system.

The filter is based upon a constant-velocity straight-line target model with position and velocity as target states. A random acceleration term was included to allow the filter to

track targets making maneuvers. The model incorporates only horizontal (east and north) position and velocity so the state evolution is described by four linear differential equations:

$$\dot{p}_E(t) = v_E(t) \quad ,$$

$$\dot{p}_E(t) = v_N(t) \quad ,$$

$$\dot{v}_E(t) = 0 \quad ,$$

$$\dot{v}_E(t) = 0 \quad .$$

Figure IV-1 illustrates the geometry of the acoustic azimuth measurement process where P is the actual target position, S is the sensor position, and P' is the position of the target at the time it emitted the sound causing the azimuth measurement. The highly nonlinear but quite simple equation for the measured azimuth, ϕ, in terms of the state variables is:

$$\phi(t) = \arctan \left\{ p_N[t - \delta(t)] - s_N,\ p_E\ [t - \delta(t)] - s_E \right\}$$

where δ is the solution of

$$c\delta(t) = \sqrt{\left\{ p_N[t - \delta(t)] - s_N \right\}^2 + \left\{ p_E\ [t - \delta(t)] - s_E \right\}^2}$$

c is the speed of sound, and s_E and s_N are the sensor positions. For subsonic targets, there is a single, closed-form solution for δ at every moment in time.

The important elements of the linearized Taylor expansion of the measurement equation are the partial derivatives of the measured azimuth in terms of the state variables:

$$\frac{\partial \phi(t)}{\partial p_E(t)} = - \left[\frac{P_N(t) - s_N}{R(t)^2} \right] \left[1 + \frac{N(t)}{D(t)} \right] \quad ,$$

$$\frac{\partial \phi(t)}{\partial p_N(t)} = \left[\frac{P_E(t) - s_E}{R(t)^2} \right] \left[1 + \frac{N(t)}{D(t)} \right] \quad ,$$

$$\frac{\partial \phi(t)}{\partial v_E(t)} = \left[\frac{P_N(t) - s_N}{CR(t)} \right] \left[\frac{1}{D(t)} \right] \quad ,$$

and

$$\frac{\partial \phi(t)}{\partial v_N(t)} = - \left[\frac{P_E(t) - s_E}{CR(t)} \right] \left[\frac{1}{D(t)} \right] \quad ,$$

where

$$N(t) = [V(t)/c] \cos [\theta(t) - \psi(t)]$$

$$D(t) = \sqrt{1 - [V(t)/c]^2 \sin [\theta(t) - \psi(t)]^2}$$

and where R, θ, V, and ψ are, respectively, the target range, the target azimuth, the target velocity, and the target heading as in Figure IV-1.

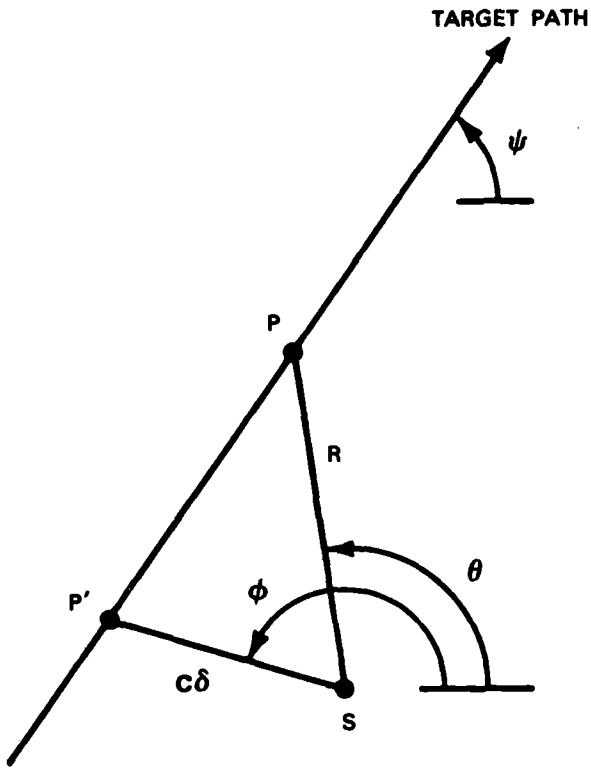

137025-N

Figure IV-1. Target/sensor geometry for acoustic azimuth measurement.

The above equations were used as the starting point for an extended Kalman filter tracker. As is typical for most highly nonlinear estimation problems, several modifications were required to handle exceptional situations, for example, when the target overflies several sensors in a row.

The filter equations were tested by a simulation equivalent to 20,000 seconds of tracker operation using azimuth measurements made every two seconds. Figure IV-2 illustrates the network of sensors for the simulation. Only 16 sensors were simulated, but a toroidally connected wraparound of target trajectories kept the target within this sensor group and effectively replicated the sensors to cover the plane. Figure IV-3 shows the probability of detection and Figure IV-4 shows the azimuth measurement accuracy use for the simulation. Target maneuvers were included in the simulation by means of a random acceleration with an rms value of 4.24 meters/second^2 that was changed every two seconds. The filter was initialized using the target's correct position and velocity and the variances were initialized to correspond to 140 meters rms position errors and 14 meters per second rms velocity errors. All azimuth measurements were centrally combined by a single filter.

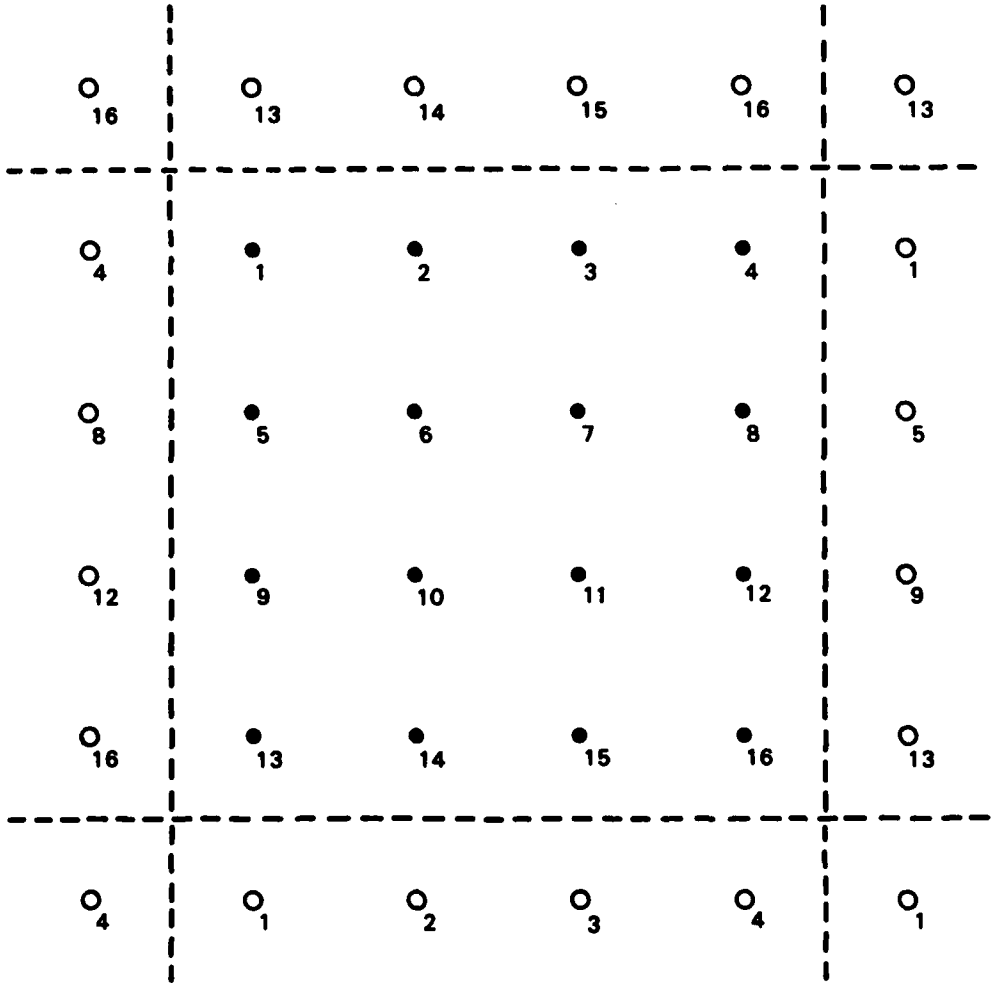

Figure IV-2. Simulated sensor array with wraparound effect indicated.

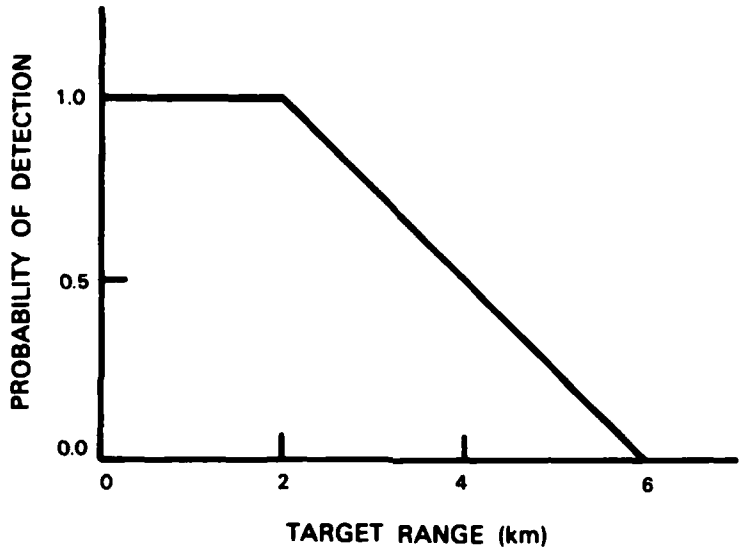

137027-N

Figure IV-3. Simulated sensor detection performance.

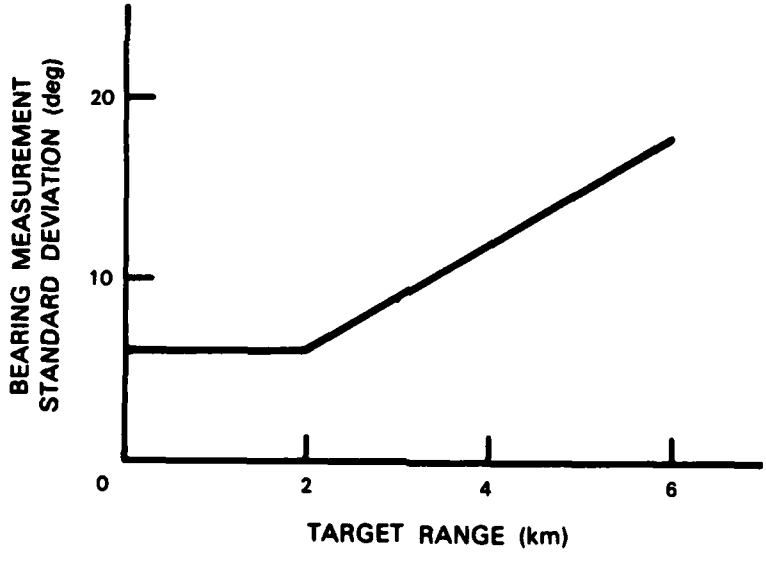

137028-N

Figure IV-4. Simulated sensor measurement performance

The simulation resulted in an average rms target position estimation error of only 227 meters and an average rms target velocity estimation error of only 9.46 meters/second. We have concluded from this that an extended Kalman filter is a practical way to implement a real-time acoustic tracker that makes direct use of azimuth measurements to update target state estimates.

B. TRACK INITIATION

The reflection method and possible position method of target location that were analyzed as reported in Section III provide target positions corresponding to some time in the past rather than corresponding to the measurement time. Pairs of measurements for two different times could be used to estimate present target locations for track initiation. However, the computational and data storage requirements as well as possible algorithmic complexity motivated an investigation of whether a more attractive alternative exists. The result has been a new and simple algorithm for obtaining real-time target location and velocity estimates for the purpose of track initiation.

Consider the situation pictured in Figure IV-5 with two nodes making measurements of a single target. P shows the true target position at the measurement time. The lines radiating from the two sensors indicate the acoustic azimuth for each node at the measurement time. We have derived the following equation to solve for P and the target velocity given the azimuth and azimuth rates at the two nodes.

$$\begin{bmatrix} \cos\phi_1(t) & -\sin\phi_1(t) & s_{1N}/c & -s_{1E}/c \\ \cos\phi_2(t) & -\sin\phi_2(t) & s_{2N}/c & -s_{2E}/c \\ \dot{\phi}_1(t)\sin\phi_{1E}(t) & \phi_1(t)\cos\phi_1 & -\cos\phi_1(t) & \sin\phi_1(t) \\ \dot{\phi}_2(t)\sin\phi_2(t) & \phi_2(t)\cos\phi_2(t) & -\cos\phi_2(t) & \sin\phi_2(t) \end{bmatrix} \begin{bmatrix} p_E(t) \\ p_N(t) \\ v_E(t) \\ v_E(t) \end{bmatrix} =$$

$$\begin{bmatrix} s_{1E}\cos\phi_1(t) - s_{1N}\sin\phi_1(t) \\ s_{2E}\cos\phi_2(t) - s_{2N}\sin\phi_2(t) \\ \dot{\phi}_1(t)s_{1E}\cos\phi_1(t) - \dot{\phi}_1(t)s_{1N}\sin\phi_1(t) \\ \dot{\phi}_2(t)s_{2E}\cos\phi_1(t) - \dot{\phi}_2(t)s_{2N}\cos\phi_2(t) \end{bmatrix} - \begin{bmatrix} 1 \\ 1 \\ 0 \\ 0 \end{bmatrix} [p_E(t)v_N(t)-p_N(t)v_E(t)],$$

where ϕ_i and $\dot{\phi}_i$ are the azimuth and azimuth rate measurements at node i and nodal locations are s_{iE} and s_{iN}. This alternative to the existing acoustic location algorithms appears attractive in that it obviates the need for accumulating azimuth histories and for searching through those histories.

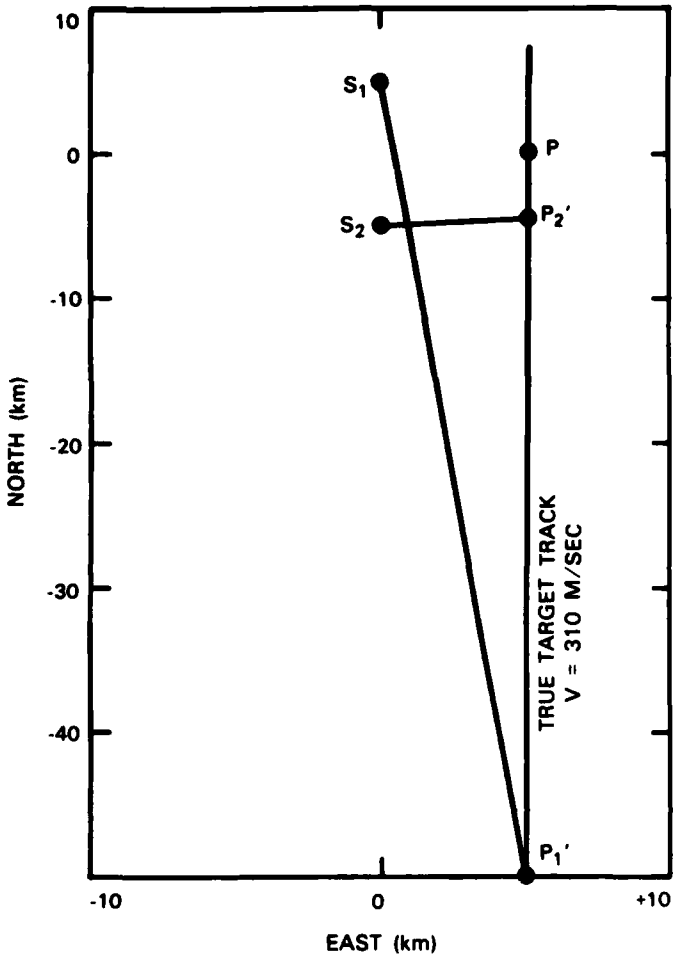

137029-N

Figure IV-5. Track initiation for target with 310 m/sec velocity and 0° heading.

In general, the solution to the above equations is not difficult and reduces to solving a quadratic equation for a single free variable to which the four state elements are linearly related. In general, there are two solutions. An example is sketched in Figure IV-6. The figure is identical to Figure IV-5, but with the extraneous trajectory added. Often, the extraneous trajectories can be immediately rejected because they involve unrealistic state variable values, but there are situations in which both solutions are physically reasonable. In those cases, it may be necessary to initiate a track with the extraneous solution and reject it only after tracking for a short time.

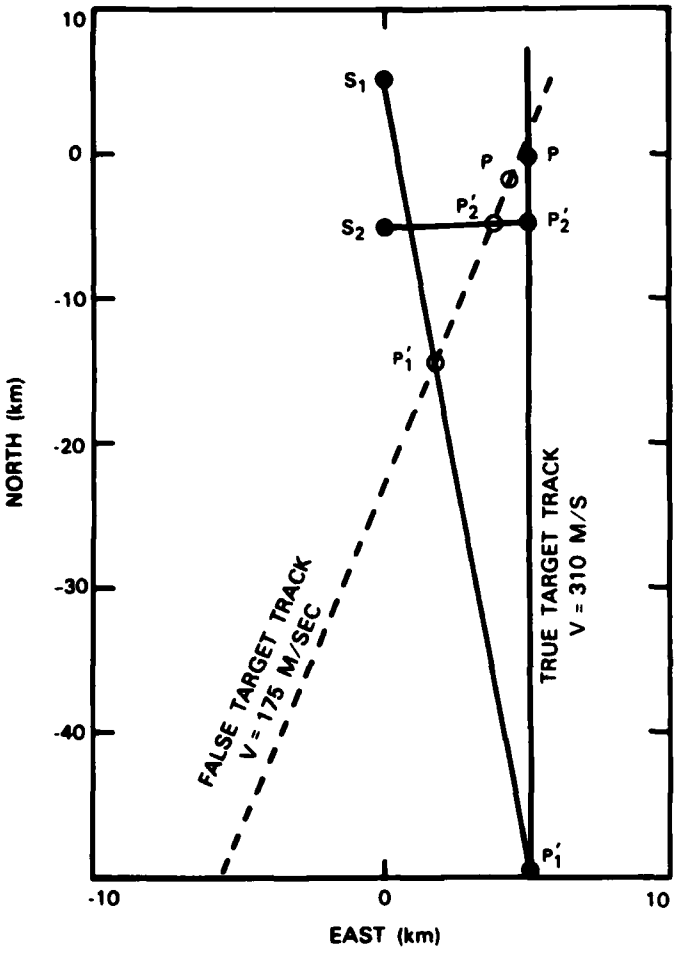

Figure IV-8. Track initiation including false target with 175 m/sec target and 10° heading.

A preliminary analysis of the sensitivity of the solutions of the above equations to azimuth tracking errors has shown that reasonable azimuth tracking errors produce acceptable errors in the estimated state elements for most target/sensor geometries. The exceptions occur when the estimated azimuths are nearly equal or nearly opposite. These geometries correspond to a target with a trajectory close to the line passing through both sensors or to a target very distant from both sensors. The previously developed acoustic location methods also perform poorly for these situations in which azimuth measurements do not provide sufficient information for estimating target positions.

C. DISTRIBUTED TRACKING

An investigation of options for distributed tracking has shown that it should be possible to develop distributed algorithms and adaptive communication policies that will approximate the tracking performance of a centralized algorithm but with reduced internodal communication. The key idea is that a target track, including error estimates, summarizes all of the past information which went into forming it and Bayesian estimation techniques can be used to optimally combine tracks. Figure IV-7 schematically shows the concept. Two nodes begin with common state variable estimates and then independently track the target using their own azimuth measurements. After a period of time, the target tracks formed by the two nodes will differ, but Bayesian techniques can be used to combine the tracks and obtain the same result as would have been obtained by continuous centralized updating with both measurement sets.

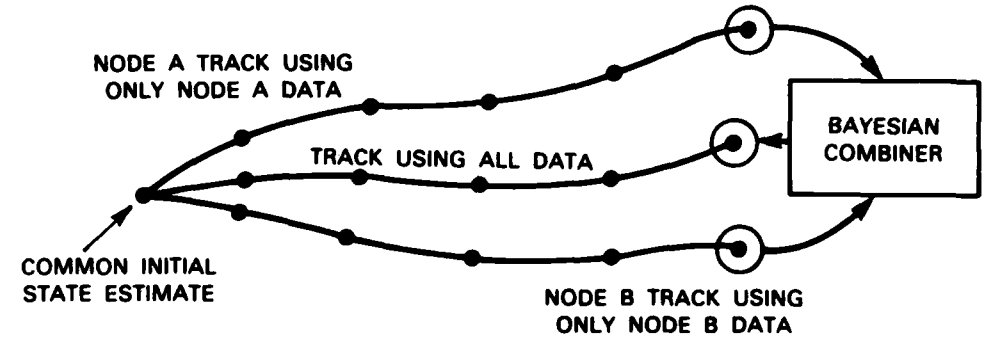

134471-N-01

Figure IV-7. Bayesian track combiner concept.

Figure IV-8 shows a nodal tracking algorithm organization that we have developed based upon these ideas. Each node maintains two tracks for each target. One is updated using local azimuth measurements (the local track) and the other is not (the common track). The common track is a Kalman prediction of future state variables for tracks received from other nodes. The local track contains that same information plus local information not yet shared with other nodes. Tracks received from outside the node are used to reset both tracks if they already exist and are used to initiate new tracks otherwise. The received tracks are used to directly reset the Kalman predictor and are processed by the Bayesian combiner to reset the local extended Kalman filter. Local azimuth measurements are used to update the local track using the extended Kalman filter. The process continues until a decision is made to broadcast the local track. When a track is broadcast, it is also used to reset the common track predictor.

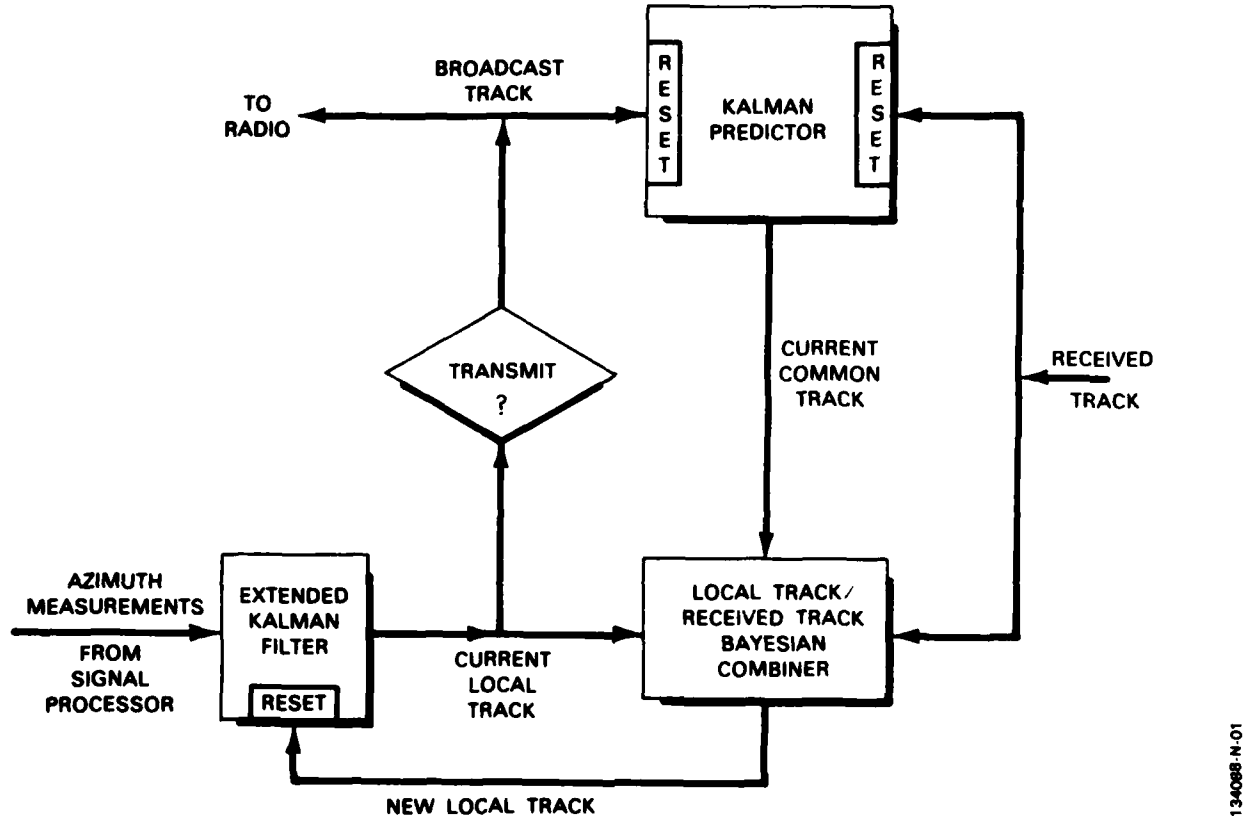

Figure IV-8. Track maintenance block diagram.

Many different policies can be formulated to decide when the nodes should broadcast their local tracks. We have experimented with one such scheme to validate our general approach to distributed tracking and adaptive communication. Figure IV-9 shows an example of the operation of this communiation policy for a simple case of a target passing through the coverage areas of three nodes. The coverage areas are shown and are assumed to be known by all nodes. The policy is for a node to broadcast its local track when it determines that the target is entering the coverage area of another node and to also broadcast when the target leaves its own coverage area. In the case of the figure, it is assumed that some other node in the network has broadcast the track received initially by node 1. This policy causes information about a target to follow the target through the DSN, temporarily residing in those nodes having the target within their coverage.

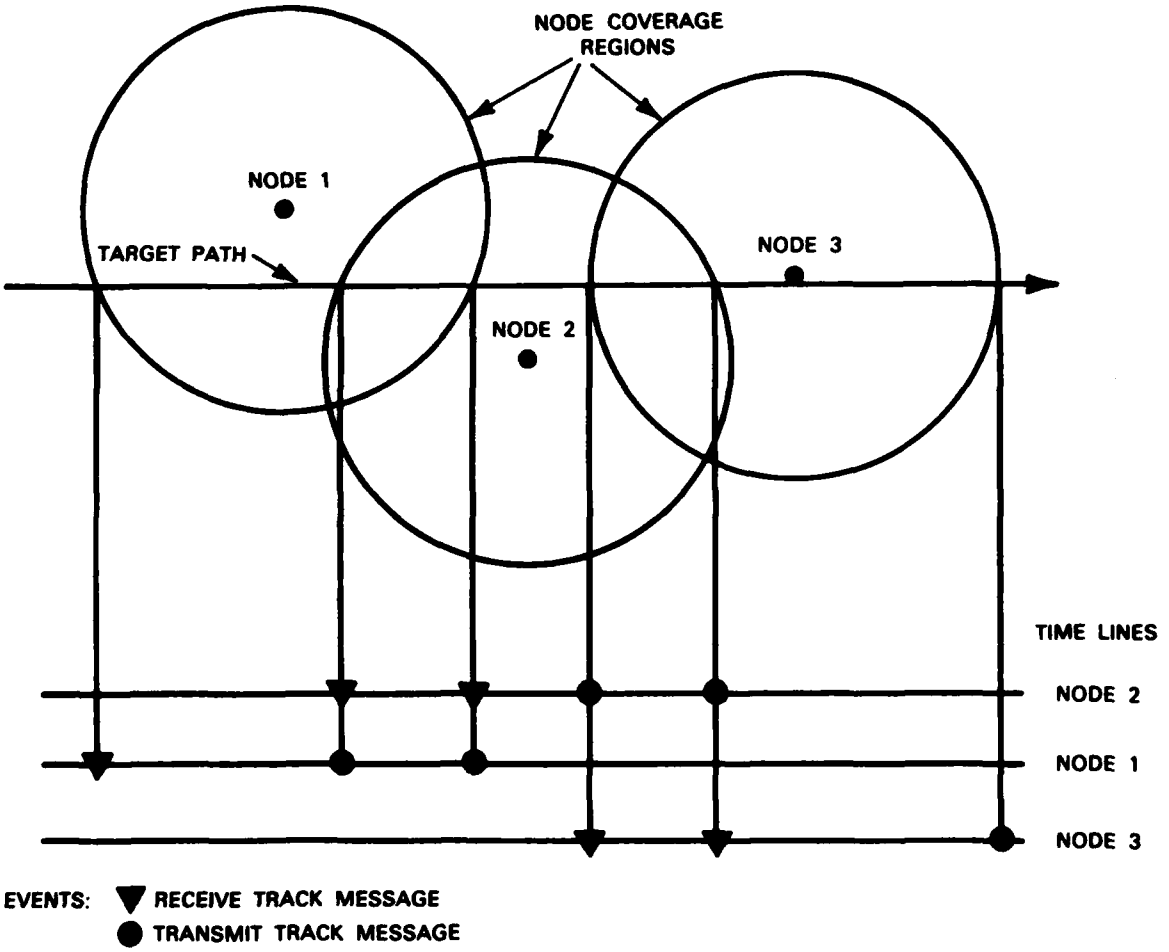

Figure IV-9. Internodal communications.

This communication policy causes tracks of the same target to differ from node to node and from the centralized extended Kalman filter tracks. Additional broadcasts would result in more consistency in target tracks among nodes and shorter periods of difference from the centralized algorithm tracks at the expense of more communications and computation.

The validity of this approach was checked by repeating the simulation of a single randomly maneuvering target passing through an indefinitely large array of sensors, but for the distributed algorithm and minimum communication policy described above. The result was an average rms target postion estimation error of 333 meters and an average rms target velocity estimation error of 9.88 meters/second, where the average is over all nodes tracking the target as well as over time. These errors are larger than in the centralized case due to the contribution of the times between broadcasts when only local tracks are updated.

D. MULTIPLE MANEUVERING TARGETS

Multiple targets increase the complication of the tracking problem. Added difficulties include the association of received track reports with local and common tracks in a node and the association of azimuth measurements with targets. These are difficult but standard problems for which many solutions have been proposed [References IV-2,3,4] that appear applicable to distributed acoustic tracking. The most promising are based upon the hypothesize-and-test paradigm popular in the artificial intelligence field. We have adapted this approach as follows. The hypotheses are sets of associations of received tracks with local and common track pairs plus associates of azimuth measurements with local tracks, some number of which are explicitly tested to see how well they fit the available data.

Maneuvering targets pose problems that are in many ways similar to the problems for multiple targets, that is, the association of tracks with tracks or measurements with tracks. Again, many solutions have been proposed [References III-3,4,5] with varying degrees of success. Some of the more recent proposals are again based on the hypothesis-and-test paradigm, attempting to solve all association-type problems in a unified framework. It is important to note that the robustness of extended Kalman filters does allow tracking of targets making moderate maneuvers without the addition of special components to the tracking algorithm.

A multiple maneuvering target simulation was carried out to confirm that our approach can be adapted to such a situation. The simulation resembled the simulations discussed in Sections IV-A and IV-C in sensor geometry, sensor performance, and wraparound of target trajectories. It resembled the Section IV-C simulation in the use of a distributed tracking algorithm. The algorithm was changed only by the addition of a simple data association component based on Reid's O-scan technique [Reference III-2]. Four targets were simulated. The targets flew 'circular' trajectories with a constant radial acceleration and a randomly perturbed tangential acceleration. The extended Kalman filter and data association algorithms were designed for constant velocity, straight-line trajectories and these random 'circular'

trajectories constituted a stressful test of their ability to associate and track multiple maneuvering targets. The simulation gave an average rms target position estimation error of 736 meters and an average rms velocity error of 25 meters/second. Given the nature of the simulated trajectories and the simplicity of the data association algorithm, this result is very encouraging.

REFERENCES

1. A. Gelb, *Applied Optimal Estimation*, M.I.T. Press, Cambridge, MA, 1974.
2. D.B. Reid, "An Algorithm for Tracking Multiple Targets", IEEE Transactions on Automatic Control, **AC-24**, No. 6., pp. 843-954 (December 1979).
3. Y. Bar-Shalom, "Tracking Methods in a Multitarget Environment", IEEE Transactions on Automatic Control, **AC-23**, No. 4., pp. 618-626 (August 1978).
4. C.Y. Chong, S. Mori, E. Tse, and R.P. Wishner, "A General Theory for Bayesian Multitarget Tracking and Classification - Generalized Tracker/Classifier (GTC)", Report TR-1015-1, Advanced Information & Decision Systems, Mountain View, CA (December 1982).
5. T. Kurien, A. Blitz, R. Washburn, and A. Willsky, "Optimal Maneuver Detection and Estimation in Multiobject Tracking", Proceedings of the 6th MIT/ONR C^3 Workshop, Cambridge, MA (July 1983).

V. APPLICATION OF ARTIFICIAL INTELLIGENCE METHODS

There are many opportunities for the application of artificial intelligence methods to the interpretation of DSN sensory data and to the control and operation of DSN systems. During this reporting period, we have carried out exploratory investigations of artificial intelligence theory and methods and have identified two DSN problem areas upon which to focus initial efforts to bring artificial intelligence methods to bear upon DSN problem areas. These two areas are the diagnosis of DSN sensory data interpretation problems experienced by DSN systems and multiple target tracking.

Expert systems protocol analysis methods* have been employed to determine the nature of expertise that is important for DSN systems, and DSN diagnosis has been selected as the area for initial expert system experimentation.

To investigate the nature of expertise in DSN systems, we used 'dialog traces' that record the kinds of communication and problem-solving that might take place while using a DSN system. This analysis clearly showed that the development of expert behavior within a DSN should not be arbitrarily focused upon traditional specialty functions such as signal processing, tracking and communications. Instead, an integrated systems viewpoint is much more amenable to expert system development. One reason is that the portion of knowledge in a specialty function area that will apply to a single system problem is usually small. Early emphasis upon individual specialty functions might result in the acquisition of specialized knowledge that would be of little use in developing higher level capabilities. The integrated system viewpoint will focus knowledge acquisition in specialty areas to only the subset required to solve the higher level problems. This observation is graphically represented in Figure V-1, which indicates that an expert system task like system diagnosis requires limited amounts of knowledge from each of several knowledge domains.

Diagnosis expertise†‡ is an area now receiving considerable attention in the field of artificial intelligence. Expert systems development tools are available that are well matched to diagnosis problems. One of these tools is the expert system program called 'EMYCIN'. We have now obtained a copy of EMYCIN and have started to use it to experiment with a rule-based approach to DSN diagnosis. The initial objectives are to evaluate the suitability

---

*F. Hayes-Roth, D.A. Waterman, D.B. Lenat, Eds., *Building Expert Systems*, Addison-Wesley Publishing Co., Reading, Mass., 1983.

†E.H. Shortliffe, *Computer-Based Medical Consultation: MYCIN*, American Elsevier, New York, 1976.

‡R. Davis, D.B. Lenat, *Knowledge-Based Systems in Artificial Intelligence*, McGraw-Hill, New York, 1980.

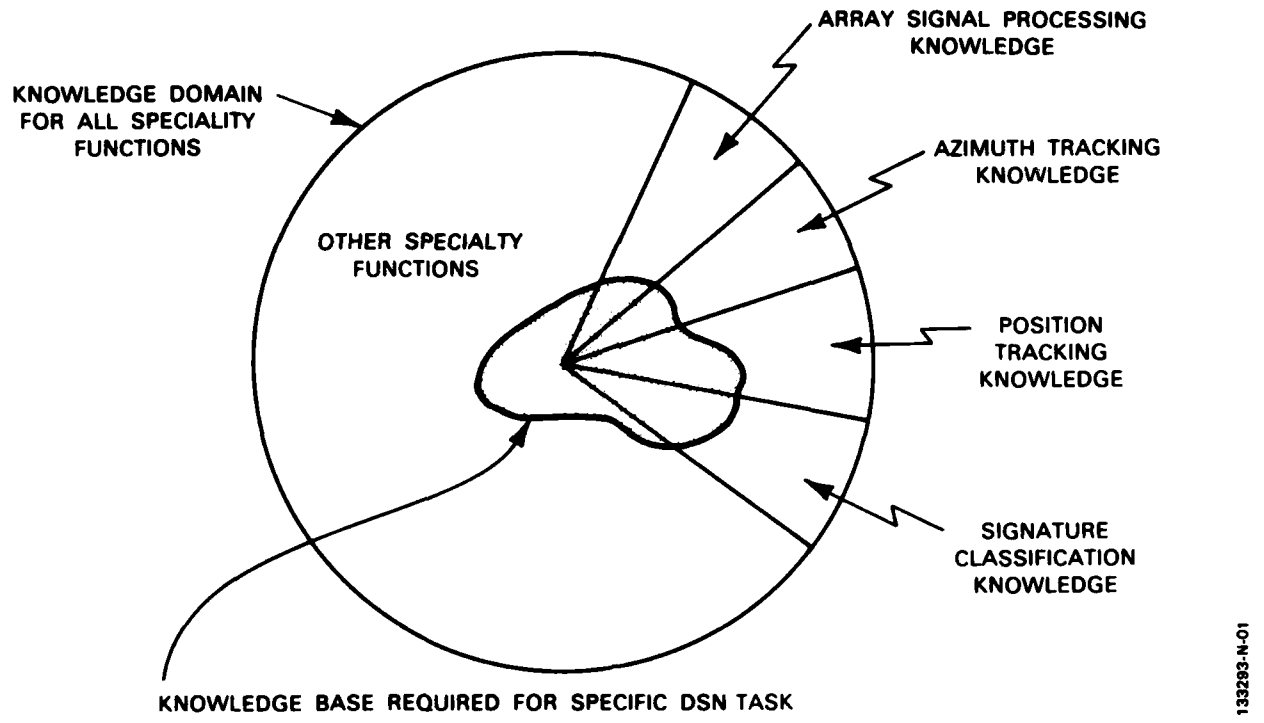

Figure V-1. Use of specialized knowledge for specific DSN problems.

of EMYCIN for developing a DSN expert diagnosis capability and to identify how its architecture should be expanded, modified or replaced. Our approach is to start with very simple programs and to implement and experiment with a series of diagnostic programs of increasing sophistication and functionality.

VI. MULTITARGET DATA-COLLECTION

Multitarget data-collection experiments were carried out at Hanscom Air Force Base during this reporting period. The major goal of these experiments was to collect DSN data for scenarios with multiple aircraft which are stressing for both signal processing and tracking functions. Issues include the ability to resolve two or more targets with close bearings as well as the ability to track a primary target in the presence of louder secondary targets. Some of this data was used for preliminary investigation of the performance of new wideband signal processing algorithms. More analysis is planned as part of the detailed evaluation of those signal processing and the new tracking algorithms described in Section IV.

A secondary goal of the data collection effort was to increase the variety of data available for evaluating DSN algorithms. These experiments were the first situations in which we used Bell-206 helicopters as targets. These aircraft are considerably smaller than the UH-1 helicopters used in previous experiments. They are quieter and present more challenging detection problems.

The route followed by targets and the locations of four DSN nodes used in the experiments are illustrated in Figure VI-1. The three different types of two-target scenarios that were executed are illustrated in Figure VI-2. These are: (1) one helicopter trailing the other with a fairly large distance between them, (2) one helicopter trailing the other with a small separation between them and (3) two helicopters flying towards and past each other at slightly different elevations. Nominal target speed was 40 meters per second (80 knots). The entire route of Figure VI-1 was traversed in approximately five minutes for each experimental run. Data recording was successful at each of the nodes for all of the experiments. The nodes labeled F and H were truck mounted nodes.

Single-target data-acquisition runs were also made using the Bell-206 helicopters. These included runs over the same tracks used for the two helicopter experiments and measurements at several different altitudes and orientations with respect to node F. These data will be used for detailed evaluation of the effect of a second target upon signal processing and tracking and to study the interactions between target elevation, target to node orientation and signal processing algorithms.

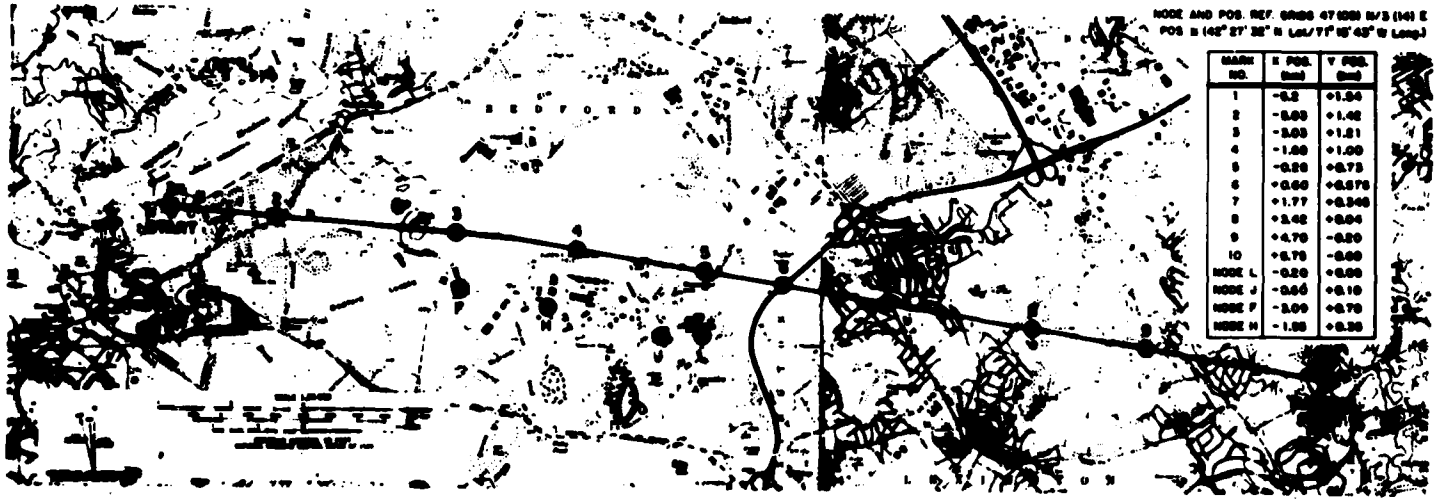

Figure VI-1. Test-bed node locations and helicopter flight paths.

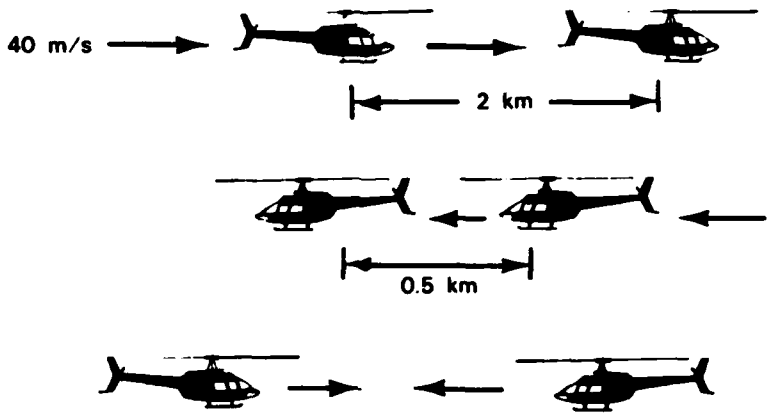

134487-N-01

Figure VI-2. Typical two-helicopter experimental scenarios.

VII. TEST BED IMPROVEMENTS

A. HARDWARE

The prototype of the new standard nodal computer described in the 31 March 1983 SATS* has been completed, tested, and integrated into the nodal signal-processing system. Figure VII-1 illustrates the new nodal configuration with the new standard nodal computer shown within the dashed lines. With the exception of the Radio Unit Interface, four additional standard nodal computers have been built and two more are under construction. Two of the additional standard nodal computer systems have also been integrated with signal-processing systems. The resulting three-node test-bed system replaces the interim system that used Versamodule single-board computers for communication and tracking functions. Integration of the remaining new nodes with signal-processing systems in three mobile nodes will be completed during the next report period. The seventh standard nodal computer will be used for maintenance and system development purposes.

The three fully integrated nodes have been operated as a three-node network to demonstrate that the hardware and all system and application software are operational. Land lines were used to provide internodal communications. A PDP-11/70 UNIX system was used to switch messages between the nodes and provide user interface and experiment-control functions. As additional standard nodal computers and signal-processing systems are integrated, they will be added to the configuration and tested in the network mode.

The test bed will shortly include a separate user-interface workstation based on a 68000-SUN processor board and the UNIX operating system. It will be possible to attach this workstation to any test-bed node. A prototype system has been procured with one megabyte of memory, a 300-Mbyte disk drive, a 9-track tape drive and a high-resolution 800- x 1000-pixel monochrome display that can be upgraded to color in the future if required.

The prototype standard node contains a radio-interface board which was described in the 31 March 1983 SATS. A second such board has now been constructed and made available to the radio-unit developers. Both boards have been tested to the extent possible without a pair of completely functional radios. Three more boards have been assembled and are being tested.

At the present time, program code must be downloaded over 9600-baud serial lines into each test-bed node from the PDP-11/70. This is a time-consuming process that interferes with optimum usage of the test bed. Floppy-disk systems have been ordered for each

*Distributed Sensor Networks Semiannual Technical Summary, Lincoln Laboratory, M.I.T. (31 March 1983) p. 27.

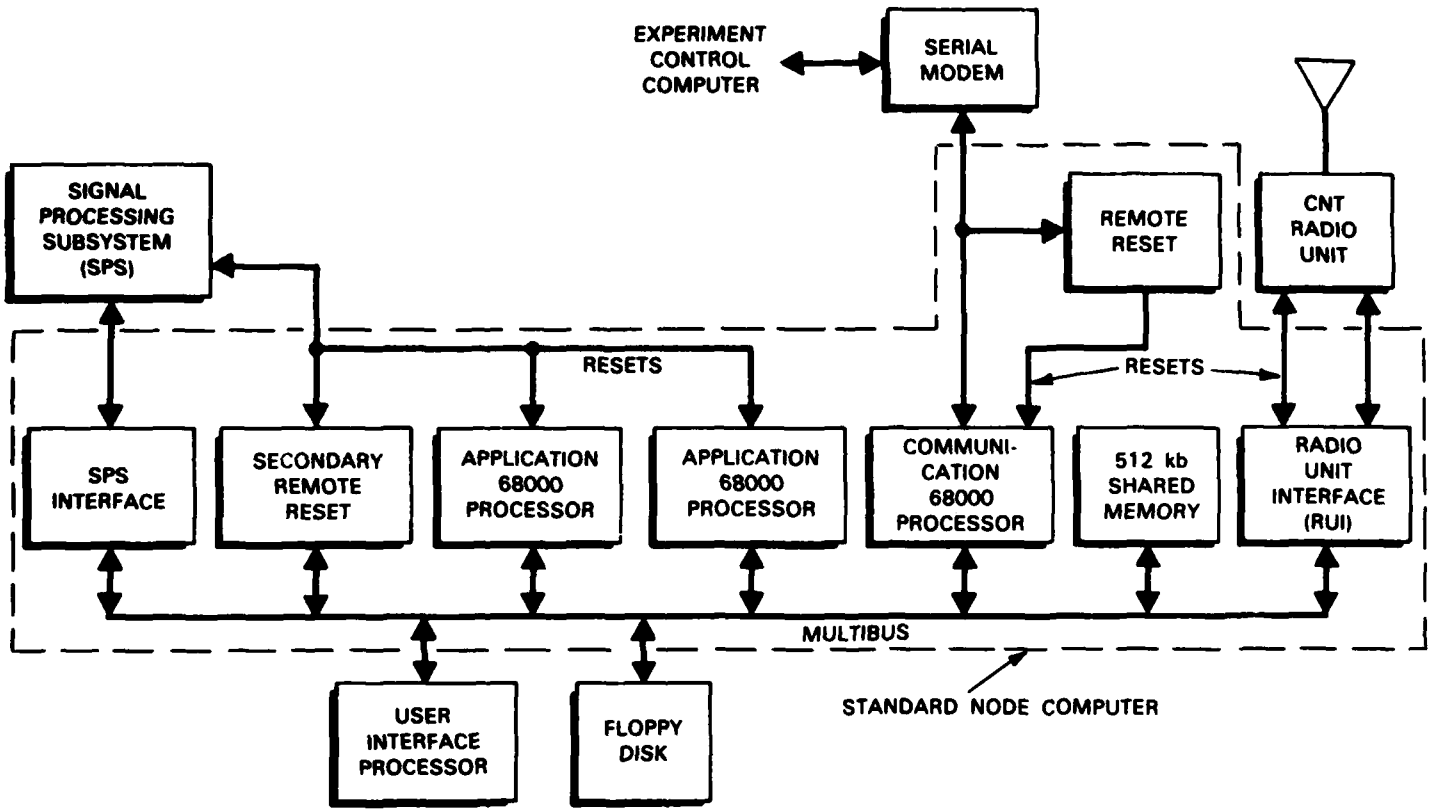

Figure VII-1. New nodal hardware configuration.

126825-N-06

test-bed node to alleviate this problem. Each unit has two double-density drives capable of storing 1.2 megabytes each. These disks will initially be used for faster loading of software into nodes and will eventually be used to record nodal data and to read simulated data into the nodes. A seventh disk system has also been ordered for our PDP-11/70 UNIX system for software development purposes.

B. SOFTWARE

Implementation of new system software and the modification and transfer of existing tracking software to the new standard nodal computer has been completed. The new system software consists of about 11,000 lines of C code and 1000 lines of assembly-language code. It occupies about 50,000 bytes of storage in the DCU SUN processor and 32,000 bytes in the applications SUN processors. The DCU storage is larger because it hosts most of the drivers. The tracking software occupies about 25,000 bytes in the applications SUN processors.

The system software includes facilities for process creation and synchronization, all necessary device servers, and a user-interface layer to perform system services such as logical device I/O. The basic nature of the system is as described in the two most recent SATS*†, although one minor and one more significant design change was made. The minor change was to make device-handling processes perform physical I/O only indirectly through the kernel. The objective was to achieve better system integrity. The more significant change was to discard messages in any outgoing broadcast message queue when the queue becomes full. This prevents internodal deadlocks in the test bed.

A series of experiments were run to measure intranodal message rates for the new standard nodal computer using the new system software. In each experiment, a process in one applications SUN processor wrote 1000 messages of a predetermined length, one immediately after the other. A process in another applications SUN processor read the 1000 messages and stopped. The time required by each processor to transfer the 1000 messages was measured. This was repeated for message lengths of 1 to 100 characters. Figure VII-2 plots the measured message rates as a function of message length. Figure VII-3 shows the message character rate (message rate times message length) versus message length. The rate is low for short messages since the software overhead must be amortized over fewer characters. As the messages become longer, the rate becomes higher because the proportional overhead is less. These rates appear satisfactory for all anticipated applications.

---

*Distributed Sensor Networks Semiannual Technical Summary, Lincoln Laboratory, M.I.T. (31 March 1983) p. 30.

†Distributed Sensor Networks Semiannual Technical Summary, Lincoln Laboratory, M.I.T. (30 September 1982) p. 29.

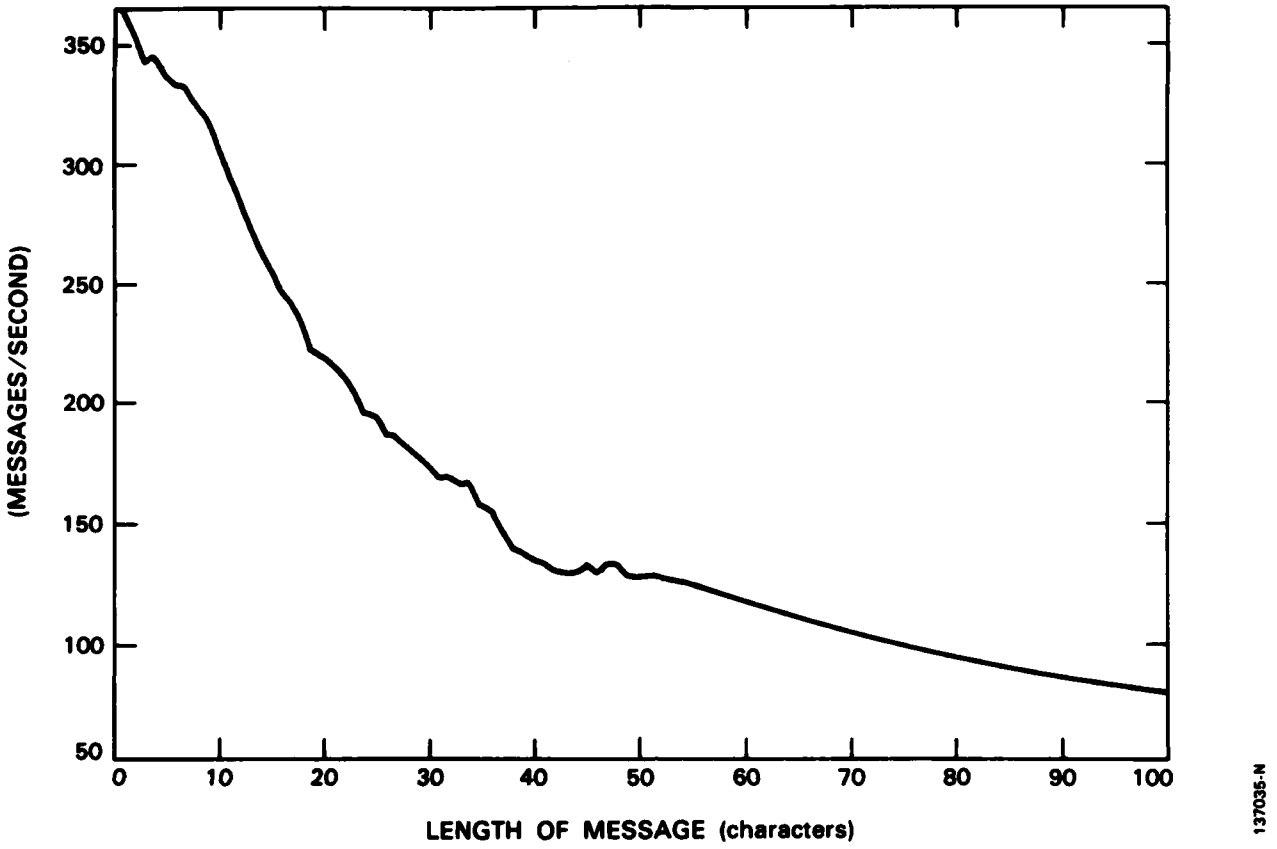

Figure VII-2. Measured intranodal message rate as a function of message length.

The system software for the new nodal configuration was designed to provide a UNIX-like interface between the system and the applications software. The objective was to allow for easy development of application software in a UNIX environment before transferring algorithms to test-bed nodes. Experience in moving the tracking software to the new nodes has proven the wisdom of this approach. Only one to two days was required for the actual transfer of the azimuth and position tracker software from the PDP-11/70 to the standard nodal computer configuration. In preparation for that transfer, a few weeks of effort was required to adapt older nodal versions of tracking software to the new UNIX interface standards.

Once the azimuth and position trackers were moved to the new nodal computers, several improvements were made. Modifications to the applications code fell into two classes: (1) changes made to reduce computational requirements of the position tracker, and (2) changes made to the user interface.

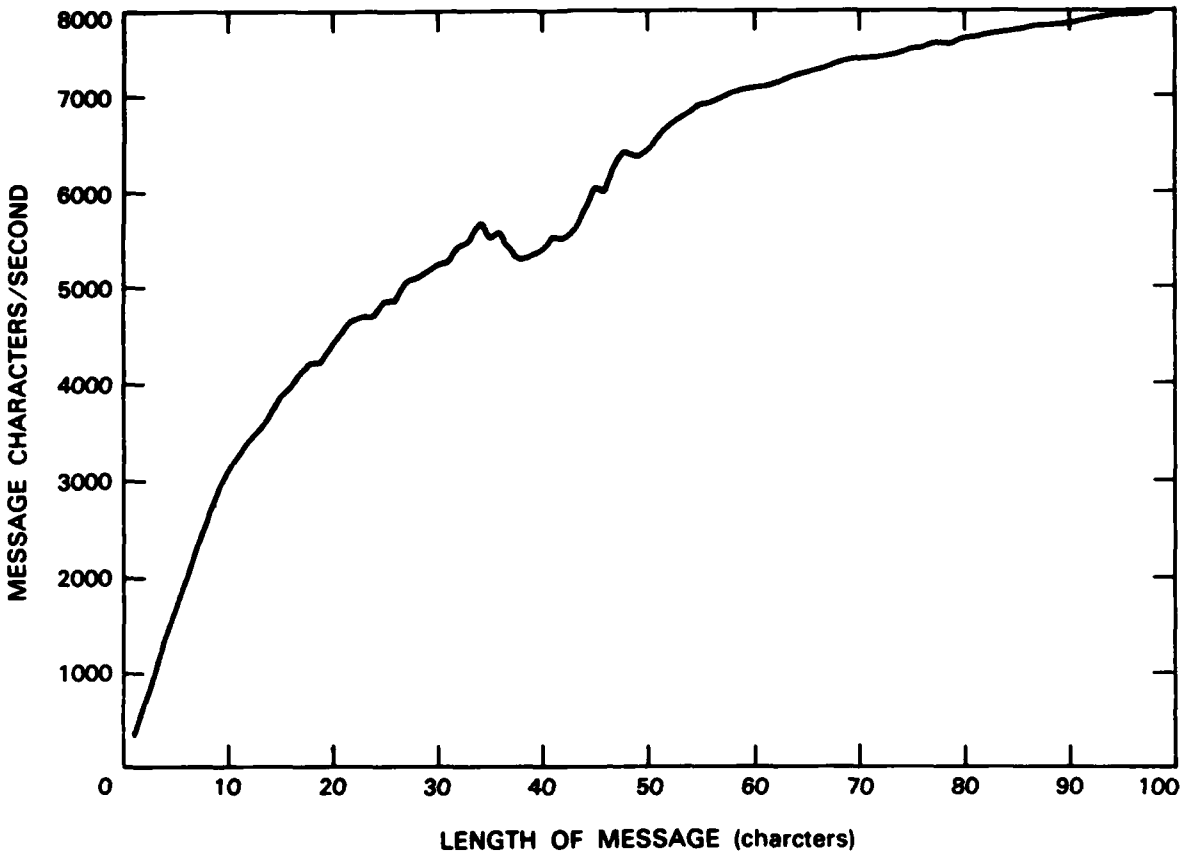

137036-N

Figure VII-3. Measured message character rate as a function of message length.

The position-tracker computational requirements were reduced by almost an order of magnitude. In the tracking experiments carried out to verify correct operation of the new nodal computers, the position trackers ran faster than real-time.

The user interface was made simpler and more flexible by introducing special interface processes that run in conjunction with each tracking process. These processes control the access to the tracking processes. They assure that commands and parameter changes for the tracking processes are complete and consistent. The interface process for each tracker is customized to that tracker. This new organization is more flexible than that previously used in that it allows applications processes to be added and changed without modifying the systems software or affecting any other application process.

An ROM-resident monitor program was provided with the SUN processor boards. It provides several critical functions including processor initialization, down-loading from a host processor, and some debugging aids. The monitor is controlled from a serial communications line. While this line can be connected to a physical terminal, it is typically connected to a

virtual terminal realized by a host computer. This monitor has thus far been used as supplied and this requires a separate serial line connection between each SUN board and a host. A more desirable situation is that all monitor functions can be realized with only a single serial line attached to a standard node computer containing multiple SUN processors. Monitor changes now have been designed and coded to provide initialization, control, remote reset and down-loading of all processors within a node over a single line. The modifications allow all processors except the DCU (to which the serial line is attached) to read commands from shared memory. The changes are now being debugged.

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*

AD-A146209

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER ESD-TR-84-002	2. GOVT ACCESSION NO. AD-A146209	3. RECIPIENT'S CATALOG NUMBER
4. TITLE *(and Subtitle)* Distributed Sensor Networks		5. TYPE OF REPORT & PERIOD COVERED Semiannual Technical Summary 1 April — 30 September 1983
		6. PERFORMING ORG. REPORT NUMBER
7. AUTHOR*(s)* Richard T. Lacoss		8. CONTRACT OR GRANT NUMBER*(s)* F19628-80-C-0002
9. PERFORMING ORGANIZATION NAME AND ADDRESS Lincoln Laboratory, M.I.T. P.O. Box 73 Lexington, MA 02173-0073		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS Program Element Nos. 61101E & 62708E Project Nos. 3D30 & 3T10 ARPA Order 3345
11. CONTROLLING OFFICE NAME AND ADDRESS Defense Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, VA 22209		12. REPORT DATE 30 September 1983 13. NUMBER OF PAGES 54
14. MONITORING AGENCY NAME & ADDRESS *(if different from Controlling Office)* Electronic Systems Division Hanscom AFB, MA 01731		15. SECURITY CLASS. *(of this report)* Unclassified 15a. DECLASSIFICATION DOWNGRADING SCHEDULE

16. DISTRIBUTION STATEMENT *(of this Report)*

Approved for public release; distribution unlimited.

17. DISTRIBUTION STATEMENT *(of the abstract entered in Block 20, if different from Report)*

18. SUPPLEMENTARY NOTES

None

19. KEY WORDS *(Continue on reverse side if necessary and identify by block number)*

multiple-sensor surveillance system
multisite detection
target surveillance and tracking
communication network
acoustic sensors
low-flying aircraft
acoustic array processing
digital radio
distributed estimation and detection

20. ABSTRACT *(Continue on reverse side if necessary and identify by block number)*

This report describes the work performed on the DARPA Distributed Sensor Networks Program at Lincoln Laboratory during the period 1 April through 30 September 1983.

DD FORM 1 Jan 73 1473 EDITION OF 1 NOV 65 IS OBSOLETE

UNCLASSIFIED

SECURITY CLASSIFICATION OF THIS PAGE *(When Data Entered)*

END

FILMED

10-84

DTIC